arXiv:1606.08707v2 [cs.FL] 15 Jul 2016

On parametric verification of asynchronous, shared-memory pushdown systems

Marie Fortin[1], Anca Muscholl[2], and Igor Walukiewicz[3]

[1]ENS Cachan, University of Paris-Saclay
[2]Technical University of Munich, IAS*
[3]University of Bordeaux, CNRS, LaBRI

Abstract

We consider the model of parametrized asynchronous shared-memory pushdown systems as introduced in [Hague'11]. In a series of recent papers it has been shown that reachability in this model is Pspace-complete [Esparza, Ganty, Majumdar'13] and that liveness is decidable in Nexptime [Durand-Gasselin, Esparza, Ganty, Majumdar'15]. We show that the liveness problem is Pspace-complete. We also consider the universal reachability problem. We show that it is decidable, and coNexptime-complete. Finally, using these results, we prove that verifying general regular properties of traces of executions, satisfying some stuttering condition, is also decidable in Nexptime for this model.

1 Introduction

It is common knowledge that even boolean programs may be impossible to analyze algorithmically. Features such as recursion or parallelism make the set of reachable configurations potentially infinite. The usual example is given by systems consisting of two pushdown processes with a shared boolean variable. Such a model can simulate a Turing machine, so every non-trivial question is undecidable [24]. Kahlon [16] proposed to consider a parametric version of this model, where the number of pushdown processes is arbitrary. At first sight this may look like a more general model, but it turns out that model-checking under various synchronization primitives is decidable, due to the lack of process identities.

Later Hague [15] considered a model where a single process has an identity, the leader process, but the operations on the shared variable do not involve any synchronization. His model of parametrized asynchronous shared-memory pushdown systems of [15] consists of one leader process and an

*On leave from University of Bordeaux, LaBRI.

arbitrary number of identical, anonymous contributor processes. Processes communicate through a shared, bounded-value register using write and read operations. A very important aspect is that there are no locks, nor test-and-set type operations, since this kind of operations would allow to elect a second leader, and all questions would be immediately undecidable. The main result of [15] is that this model still enjoys a decidable reachability problem. The complexity of this problem has been later established in [12]. Recently, Durand-Gasselin et al. [10] have also shown decidability of the liveness problem for this model.

The reachability problem in Hague's model can be formulated as whether there is a computation of the system where the leader can execute a special action, say $\top$. The *repeated reachability* problem asks if there is a computation where the leader can execute $\top$ infinitely often. This problem provides a succinct way of talking about liveness properties concerning the leader process. We also consider in this paper *universal reachability*: this is the question of deciding if on every maximal trace of the system, the leader executes $\top$. In terms of temporal logics, reachability is about EF properties, while universal reachability is about AF properties.

Our first result shows that there is no complexity gap between verification of reachability and repeated reachability in the parametrized setting, both problems are Pspace-complete. This answers the question left open by [10], that provided a Pspace lower bound and a Nexptime upper bound for the liveness problem. Technically, our Pspace upper bound requires to combine the techniques from [10] and [21]. We use a result from [10] saying that if there is a run then there is an ultimately periodic one. Then we extend the approach from [21] from finite to ultimately periodic runs.

As a second result we show that universal reachability is coNexptime-complete. This result also bears some interesting technical aspects. For the upper bound, as in the case of reachability, we use a variant of the so-called accumulator semantics [21, 10]. We need to adapt it though in order to make it sensitive to divergence. The lower bound shows that it is actually possible to force a fair amount of synchronization in the model; we can ensure that the first 2^n values written into the shared register are read in the correct order and none of them is skipped. So the coNexptime-hardness result can be interpreted positively, as showing what can be implemented in this model.

Finally, we consider properties that refer not only to the leader process, but also to contributors. As noticed in [10] such properties are undecidable in general, because they can enforce special interleavings that amount to identify a particular contributor. In the parametrized setting it is more natural to use properties that are stutter-invariant w.r.t. contributor actions. We show that parametrized verification of such properties is decidable, and establish precise complexities both when the property is given as a Büchi automaton, and as LTL formula.

Related work. Parametrized verification of shared-memory, multi-threaded programs has been studied for finite-state threads e.g. in [2, 17] and for pushdown threads in [1, 5, 19, 20]. The decidability results in [1, 5, 19, 20] concern the reachability analysis up to a bounded number of execution contexts, and in [1, 5], dynamic thread creation is allowed. The main difference with our setting is that synchronization primitives are allowed in those models, so decidability depends on restricting the set of executions in the spirit of bounded context switches. Our model does not have such a restriction, but forbids synchronization instead.

Besides the already cited papers, a related paper that goes beyond the reachability property is [6]. Bouyer et. al. consider in [6] a very similar setting, but without leader and only finite-state contributors. They consider the problem of *almost-sure* reachability, which asks if a given state is reached by some process with probability 1 under a stochastic scheduler. They exhibit the existence of positive or negative cut-offs, and show that the problem can be decided in EXPSPACE, and is at least PSPACE-hard. By contrast, the universal reachability problem we consider here, although at first glance close to the question in [6], has very different characteristics. It is in NP for finite-state contributors, and can be simply solved by using a variant of the accumulator semantics. The challenge in the present paper comes from considering pushdown systems.

Finally, we should mention that there is a rich literature concerning the verification of *asynchronously-communicating* parametrized programs, that is mostly related to the verification of distributed protocols and varies for approaches and models (see e.g. [14, 7, 11, 18] for some early work, and [9, 22, 3] and references therein). Most of these papers are concerned with finite-state programs only, which are not the main focus of our results.

Outline of the paper. In Section 2, we introduce Hague's model. Section 3 presents the problems considered in this paper, and gives an overview of our results. In Section 4 we study the repeated reachability problem, and in Section 5 universal reachability. Finally, we show in Sections 6 and 7 how these results may be used to verify more general properties.

2 Preliminaries

In this section we recall the model of parametrized systems of [15], which we call $(\mathcal{C}, \mathcal{D})$-systems. First we give some basic definitions and notations.

2.1 Standard definitions

A *multiset* over a set E is a function $M : E \to \mathbb{N}$. We let $|M| = \sum_{x \in E} M(x)$. The *support* of M is the set $\{x \in E \mid M(x) > 0\}$. For $n \in \mathbb{N}$, we write nM, $M + M'$ and $M - M'$ for the multisets defined by $(nM)(x) = n \cdot M(x)$,

$(M + M')(x) = M(x) + M'(x)$ and $(M - M')(x) = \max(0, M(x) - M'(x))$. We denote by $[x]$ the multiset containing a single copy of x, and $[x_1, \ldots, x_n]$ the multiset $[x_1] + \ldots + [x_n]$. We write $M \leq M'$ when $M(x) \leq M'(x)$ for all x.

A *transition system* over a finite alphabet Σ is a tuple $\langle S, \delta, s_{init} \rangle$ where S is a (finite or infinite) set of states, $\delta \subseteq S \times \Sigma \times S$ is a set of transitions, and $s_{init} \in S$ the initial state. We write $s \xrightarrow{u} s'$ (for $u \in \Sigma^*$) when there exists a path from s to s' labeled by u. A *trace* is a sequence of actions labeling a path starting in s_{init}; so u is a trace if $s_{init} \xrightarrow{u} s'$ for some s'.

A *pushdown system* is a tuple $\langle Q, \Sigma, \Gamma, \Delta, q_{init}, A_{init} \rangle$ consisting of a finite set of states Q, a finite input alphabet Σ, a finite stack alphabet Γ, a set of transitions $\Delta \subseteq (Q \times \Gamma) \times (\Sigma \cup \{\varepsilon\}) \times (Q \times \Gamma^*)$, an initial state $q_{init} \in Q$, and an initial stack symbol $A_{init} \in \Gamma$. The associated transition system has $Q \times \Gamma^*$ as states, $q_{init} A_{init}$ as the initial state, and transitions $qA\alpha \xrightarrow{a} q'\alpha'\alpha$ for $(q, A, a, q', \alpha') \in \Delta$.

A word $u = a_1 \cdots a_n$ is a *subword* of v (written $u \sqsubseteq v$) when there are words $v_0, \ldots, v_n$ such that $v = v_0 a_1 v_1 \cdots v_{n-1} a_n v_n$, so u is obtained from v by erasing symbols. The *downward closure* of a language $L \subseteq \Sigma^*$ is $L{\downarrow} = \{u \in \Sigma^* \mid \exists v \in L.\, u \sqsubseteq v\}$.

2.2 $(\mathcal{C}, \mathcal{D})$-systems

We proceed to the formal definition of $(\mathcal{C}, \mathcal{D})$-systems. These systems are composed of arbitrary many instances of a contributor process $\mathcal{C}$ and one instance of a leader process $\mathcal{D}$. The processes communicate through a shared register. We write G for the finite set of register values, and use g, h to range over elements of G. The initial value of the register is denoted by g_{init}. The alphabets of both $\mathcal{C}$ and $\mathcal{D}$ contain actions representing reads and writes to the register:

$$\Sigma_C = \{\overline{r}(g), \overline{w}(g) : g \in G\} \qquad \Sigma_D = \{r(g), w(g) : g \in G\}\,.$$

Both $\mathcal{C}$ and $\mathcal{D}$ are (possibly infinite) transition systems over these alphabets:

$$\mathcal{C} = \langle S, \delta \subseteq S \times \Sigma_C \times S, s_{init} \rangle \qquad \mathcal{D} = \langle T, \Delta \subseteq T \times \Sigma_D \times T, t_{init} \rangle \qquad (1)$$

In this paper we will be interested in the special case where $\mathcal{C}$ and $\mathcal{D}$ are pushdown transition systems:

$$\mathcal{A}_C = \langle P, \Sigma_C, \Gamma_C, \delta, p_{init}, A^C_{init} \rangle \qquad \mathcal{A}_D = \langle Q, \Sigma_D, \Gamma_D, \Delta, q_{init}, A^D_{init} \rangle \qquad (2)$$

In this case the transition system $\mathcal{C}$ from (1) is the transition system associated with $\mathcal{A}_C$: its set of states is $S = P \times (\Gamma_C)^*$ and the transition relation δ is defined by the push and pop operations. Similarly, the transition system $\mathcal{D}$ is determined by $\mathcal{A}_D$. When stating general results on $(\mathcal{C}, \mathcal{D})$-systems we will use the notations from Eq. (1); when we need to refer to precise states, or

use some particular property of pushdown transition systems, we will employ the notations from Eq. (2).

A $(\mathcal{C},\mathcal{D})$-system consists of an arbitrary number of copies of $\mathcal{C}$, one copy of $\mathcal{D}$, and a shared register. So a *configuration* is a triple $(M \in \mathbb{N}^S, t \in T, g \in G)$, consisting of a multiset M counting the number of instances of $\mathcal{C}$ in a given state, the state t of $\mathcal{D}$ and the current register value g.

In order to define the transitions of the $(\mathcal{C},\mathcal{D})$-system we extend the transition relation δ of $\mathcal{C}$ from elements of S to multisets over S:

$$M \xrightarrow{a} M' \text{ in } \delta \qquad \text{if } s \xrightarrow{a} s' \text{ in } \delta,\ M(s) > 0, \text{ and } M' = M - [s] + [s'], \text{ for some } s, s' \in S.$$

Observe that such a transition does not change the size of the multiset. The transitions of the $(\mathcal{C},\mathcal{D})$-system are either transitions of the leader (the first two cases below) or transitions of contributors (last two cases):

$$\begin{aligned}
(M,t,g) &\xrightarrow{w(h)} (M,t',h) && \text{if } t \xrightarrow{w(h)} t' \text{ in } \Delta\,,\\
(M,t,g) &\xrightarrow{r(h)} (M,t',h) && \text{if } t \xrightarrow{r(h)} t' \text{ in } \Delta \text{ and } h = g\,,\\
(M,t,g) &\xrightarrow{\overline{w}(h)} (M',t,h) && \text{if } M \xrightarrow{\overline{w}(h)} M' \text{ in } \delta\,,\\
(M,t,g) &\xrightarrow{\overline{r}(h)} (M',t,h) && \text{if } M \xrightarrow{\overline{r}(h)} M' \text{ in } \delta \text{ and } h = g\,.
\end{aligned}$$

A *run* from a configuration (M,t,g) is a finite or an infinite sequence of transitions starting in (M,t,g). A run can start with any number n of contributors, but then the number of contributors is constant during the run. A run is *initial* if it starts in a configuration of the form $(n[s_{init}], t_{init}, g_{init})$, for some $n \in \mathbb{N}$. It is *maximal* if it is initial and is not a prefix of any other run. In particular, every infinite initial run is maximal. A *(maximal) trace* of the $(\mathcal{C},\mathcal{D})$-system is a finite or an infinite sequence over $\Sigma_C \cup \Sigma_D$ labeling a (maximal) initial run.

3 Problem statement and overview of results

We are interested in verifying properties of traces of $(\mathcal{C},\mathcal{D})$-systems. The first question is what kind of specifications we may use. In general, verification of regular, action-based properties of runs of pushdown $(\mathcal{C},\mathcal{D})$-systems is undecidable: such properties allow to control the interleavings of contributors and thus identify e.g. a single contributor that runs together with the leader (see e.g. [10]).

For this reason we consider $\mathcal{C}$*-expanding* properties $\mathcal{P} \subseteq (\Sigma_C \cup \Sigma_D)^\infty$. By this we mean properties where actions of contributors can be replicated: if $u_0 a_0 u_1 a_1 u_2 \cdots \in \mathcal{P}$ with $a_i \in \Sigma_C$, $u_i \in \Sigma_D^*$, and $f : \mathbb{N} \to \mathbb{N}^+$, then

$u_0 a_0^{f(0)} u_1 a_1^{f(1)} u_2 \cdots \in \mathcal{P}$, too. Because of parametrization, contributor actions can always be replicated, so $\mathcal{C}$-expanding properties are a natural class of properties for $(\mathcal{C}, \mathcal{D})$-systems.

A related, more classical notion is *stutter-invariance*. A language $L \subseteq \Sigma^\infty$ is stutter-invariant if for every finite or infinite sequence $a_0 a_1 \cdots$ and every function $f : \mathbb{N} \to \mathbb{N}^+$, we have $a_0 a_1 \cdots \in L$ iff $a_0^{f(0)} a_1^{f(1)} \cdots \in L$. It is known that the stutter-invariant properties expressible in linear-time temporal logic LTL are precisely those expressible in LTL without the next-operator [23, 13]. By definition, every stutter-invariant property is $\mathcal{C}$-expanding.

We will consider regular properties $\mathcal{P} \subseteq (\Sigma_C \cup \Sigma_D)^\infty$ that are $\mathcal{C}$-expanding, as defined above. These properties will be described either by an LTL formula, or by an automaton $\mathcal{A} = \langle Q, \Sigma_C \cup \Sigma_D, \Delta, q_0, F, R \rangle$ with finite set of states Q, alphabet $\Sigma_C \cup \Sigma_D$, transitions $\Delta \subseteq Q \times (\Sigma_C \cup \Sigma_D) \times Q$, and sets of final states $F, R \subseteq Q$. Finite runs are accepted if they end in F, infinite ones are accepted if they visit R infinitely often. For simplicity, we call such an automaton a *Büchi automaton*.

We will also consider some particular properties, that are essential for the general decision procedure:

1. The *reachability problem* asks if the $(\mathcal{C}, \mathcal{D})$-system has some trace containing a given leader action $\top$.

2. The *repeated reachability problem* asks if the $(\mathcal{C}, \mathcal{D})$-system has some trace with infinitely many occurrences of a given leader action $\top$.

3. The *universal reachability problem* asks if every maximal trace of the $(\mathcal{C}, \mathcal{D})$-system contains a given leader action $\top$.

4. The complement of the previous question is the *max-safe problem*. It asks if the $(\mathcal{C}, \mathcal{D})$-system has some maximal trace that does not contain a given leader action $\top$.

The above problems are of course basic examples of (stutter-invariant) LTL properties over $\Sigma_D \cup \Sigma_C$. In the following example we show how they can be used together to verify some more involved properties of parametrized systems.

Example 1 The consensus problem consists in making all processes agree on a common value, among those values that were proposed by the processes. For simplicity we can assume that each (leader or contributor) process selects initially some value $b \in \{0, 1\}$, and has two special actions, $choose(0)$ and $choose(1)$. If a (leader or contributor) process performs $choose(i)$, this means that a value has been agreed upon, and it is equal to i. Furthermore, we denote the set of actions that are possible after $choose(i)$ as Σ_i, and assume that $\Sigma_0 \cap \Sigma_1 = \emptyset$.

The property we are interested in asks that processes should agree on a common value from $\{0,1\}$, and then use this value in all future computations. Stated as a property of the leader and contributors it means:

$$\mathsf{AF}(\bigvee_{b=0,1} \mathit{choose}(b)) \wedge \mathsf{AG}(\bigwedge_{b=0,1} (\mathit{choose}(b) \Rightarrow \mathsf{AG}\,\Sigma_b))$$

The first part of the property corresponds to universal reachability. The second one is a safety property (the complement of a reachability property), requiring that there is no maximal run containing action $\mathit{choose}(b)$ and later on, some action from Σ_{1-b}.

The main results of our paper establish the precise complexity of all questions about $(\mathcal{C},\mathcal{D})$-systems that were introduced above. We first state the result concerning the largest class of properties, and then for specific cases. The proofs follow the inverse order, the results about specific cases are used to prove the general result.

Theorem 2 *The following problem is* NEXPTIME*-complete: given a pushdown $(\mathcal{C},\mathcal{D})$-system and a $\mathcal{C}$-expanding regular property $\mathcal{P}$ over $\Sigma_D \cup \Sigma_C$ (given by a Büchi automaton or by an LTL formula), determine if the $(\mathcal{C},\mathcal{D})$-system has a maximal trace in $\mathcal{P}$.*

In principle, it could be more difficult to verify properties given by LTL formulas since LTL formulas can be exponentially more succinct than non-deterministic Büchi automata. The above theorem implies that this blowup in the translation from LTL to non-deterministic Büchi automata does not influence the complexity of the algorithm. In this context, it is worth to recall that even for a single pushdown automaton, LTL model-checking is EXPTIME-complete [4].

Theorem 3 *The repeated reachability problem for pushdown $(\mathcal{C},\mathcal{D})$-systems is* PSPACE*-complete.*

Theorem 4 *The max-safe problem for pushdown $(\mathcal{C},\mathcal{D})$-systems is* NEXPTIME*-complete. It is* NP*-complete when $\mathcal{C}$ ranges over finite-state systems.*

The proof of Theorem 2 uses Theorems 3 and 4 and one more result, that is interesting on its own. Note that both the repeated reachability and the max-safe problem talk about one distinguished action of the leader, while $\mathcal{C}$-expanding properties also refer to actions of contributors. Perhaps a bit surprisingly, we show how to modify a $(\mathcal{C},\mathcal{D})$-system so that only leader actions matter. We reduce the problem of verifying if a $(\mathcal{C},\mathcal{D})$-system satisfies a $\mathcal{C}$-expanding property $\mathcal{P}$ to the problem of verifying that some polynomially larger $(\tilde{\mathcal{C}},\tilde{\mathcal{D}})$-system satisfies a property $\tilde{\mathcal{P}}$ that refers *only* to actions of the leader $\tilde{\mathcal{D}}$. The idea is that the leader is given the control of the register, and

contributors just submit read or write requests; these are processed by the leader who later sends acknowledgements to the contributors. The result of this reduction is summarized by the following theorem:

Theorem 5 *For every $(\mathcal{C},\mathcal{D})$-system, there exists a $(\tilde{\mathcal{C}},\tilde{\mathcal{D}})$-system such that for every $\mathcal{C}$-expanding property $\mathcal{P} \subseteq (\Sigma_D \cup \Sigma_C)^\infty$, there exists a property $\tilde{\mathcal{P}} \subseteq (\widetilde{\Sigma}_D)^\infty$, where $\widetilde{\Sigma}_D$ is the action alphabet of $\tilde{\mathcal{D}}$, such that:*

1. *the $(\mathcal{C},\mathcal{D})$-system has a finite (resp. infinite) maximal trace in $\mathcal{P}$ iff the $(\tilde{\mathcal{C}},\tilde{\mathcal{D}})$-system has a finite (resp. infinite) maximal trace whose projection on $\widetilde{\Sigma}_D$ is in $\tilde{\mathcal{P}}$;*
2. *every infinite run of the $(\tilde{\mathcal{C}},\tilde{\mathcal{D}})$-system has infinitely many write operations of $\tilde{\mathcal{D}}$;*
3. *the $(\tilde{\mathcal{C}},\tilde{\mathcal{D}})$-system has an infinite run iff the $(\mathcal{C},\mathcal{D})$-system has one.*

If the $(\mathcal{C},\mathcal{D})$-system is defined by pushdown automata $\mathcal{A}_C$ and $\mathcal{A}_D$, then the $(\tilde{\mathcal{C}},\tilde{\mathcal{D}})$-system is effectively defined by pushdown automata of sizes linear in the sizes of $\mathcal{A}_C$ and $\mathcal{A}_D$. If $\mathcal{P}$ is a regular, respectively LTL property, then so is $\tilde{\mathcal{P}}$. An automaton or LTL formula of linear size for $\tilde{\mathcal{P}}$ is effectively computable from the one for $\mathcal{P}$.

In the remaining of the paper we successively give the proofs of Theorem 3, Theorem 4, Theorem 2, and Theorem 5. Though it is proven in the last section, Theorem 5 does not rely on other results, and will be used in the proofs of Theorem 4 and 2.

4 Repeated reachability

We show in this section that repeated reachability for pushdown $(\mathcal{C},\mathcal{D})$-systems can be decided in PSPACE (Theorem 3). The matching lower bound comes from the PSPACE lower bound for the reachability problem [12]. We call a run of the $(\mathcal{C},\mathcal{D})$-system a *Büchi run* if it has infinitely many occurrences of the leader action $\top$. So the problem is to decide if a given $(\mathcal{C},\mathcal{D})$-system has a Büchi run.

Our proof has three steps. The first one relies on a result from [10], showing that the stacks of contributors can be assumed to be polynomially bounded. This allows to search for ultimately periodic runs (Lemma 9), as in the case of one pushdown system. The next step extends the capacity technique introduced in [21] for the reachability problem, to Büchi runs. We reduce the search for Büchi runs to the existence of ω-supported runs (Lemma 11). The last step is the observation that, as in the case of finite runs, we can use the downward closure of capacity runs of the leader (Lemma 12). Overall this yields a PSPACE algorithm for the existence of Büchi runs (Theorem 3).

4.1 Finite-state contributors

As observed in [12], pushdown contributors can be simulated by finite-state ones, by exploiting the fact that the setting is parametrized. To state the result we need the notion of *effective stack-height* for a pushdown system. *Effective* stack-height refers to the part of the stack that is still used in the future. Consider a possibly infinite run $\rho = q_1\alpha_1 \xrightarrow{a_1} q_2\alpha_2 \xrightarrow{a_2} \dots$ of a pushdown system. We write $\alpha_i = \alpha'_i\alpha''_i$, where α''_i is the longest suffix of α_i that is also a proper suffix of α_j for all $j > i$. The *effective stack-height* of a configuration $q_i\alpha_i$ in ρ is the length of α'_i. (Notice that even though it is never popped, the first element of the longest common suffix of the $(\alpha_i)_{j\geq i}$ may be read, hence the use of *proper* suffixes.)

By $\mathcal{C}_N$ we denote the restriction of the contributor pushdown $\mathcal{A}_C$ to runs in which all configurations have effective stack-height at most N, where N is a positive integer. More precisely, $\mathcal{C}_N$ is the finite-state system with set of states $\{p\alpha \in P\Gamma_C^* : |\alpha| \leq N\}$, and transitions $p\alpha \xrightarrow{a} q\alpha'$ if $p\alpha \xrightarrow{a} q\alpha'\alpha''$ in Δ for some α''. Note that $\mathcal{C}_N$ is effectively computable in PSPACE from $\mathcal{A}_C$ and N given in unary. The key idea in [10] is that when looking for Büchi runs for pushdown $(\mathcal{C},\mathcal{D})$-systems, $\mathcal{C}$ can be replaced by $\mathcal{C}_N$ for N polynomially bounded:

Theorem 6 (Thm. 4 in [10]) *Let $N > 2|P|^2|\Gamma_C|$. There is a Büchi run in the $(\mathcal{C},\mathcal{D})$-system iff there is one in the $(\mathcal{C}_N,\mathcal{D})$-system.*

A similar result for finite runs can be derived from the proof of this theorem.

Lemma 7 Let $N > 2|P|^2|\Gamma_C| + 1$. A configuration $([p_1\alpha_1,\dots,p_n\alpha_n],t,g)$ of the $(\mathcal{C}_N,\mathcal{D})$-system is reachable iff there exists a reachable configuration of the $(\mathcal{C},\mathcal{D})$-system of the form $([p_1\alpha_1\beta_1,\dots,p_n\alpha_n\beta_n],t,g)$, for some β_i.

Notation. Throughout the paper, we write $\mathcal{C}_{fin}$ for $\mathcal{C}_N$ with $N = 2|P|^2|\Gamma_C| + 2$. We will use the notation $\langle P_{fin},\Sigma_C,\delta,p^{fin}_{init}\rangle$ for the finite-state system $\mathcal{C}_{fin}$, and continue to write $\mathcal{A}_D = \langle Q,\Sigma_D,\Gamma_D,\Delta,q_{init},A^D_{init}\rangle$ for the pushdown system $\mathcal{D}$.

Remark 8 It is not difficult to see that every infinite run of a pushdown system contains infinitely many configurations with effective stack-height one (see also [10]).

Putting together Theorem 6 and Remark 8, we obtain:

Lemma 9 There is a Büchi run in the $(\mathcal{C}_{fin},\mathcal{D})$-system iff there is one of the form

$$(n[p^{fin}_{init}],t_{init},g_{init}) \xrightarrow{u} (M,t_1,g) \xrightarrow{v} (M,t_2,g) \xrightarrow{v} \dots$$

for some $n \in \mathbb{N}$, $g \in G$, $M \in (P_{fin})^n$, $u,v \in (\Sigma_C \cup \Sigma_D)^*$, where:

- v ends by a letter from Σ_D and contains $\top$, and
- all configurations $t_i \in Q\Gamma_D^*$ of $\mathcal{D}$ have effective stack-height one, the same control state, and the same top stack symbol.

Proof. Let $\rho = (n[p_{init}^{fin}], t_{init}, g_{init}) = (M_0, t_0, g_0) \xrightarrow{a_1} (M_1, t_1, g_1) \xrightarrow{a_2} \cdots$ be a Büchi run of the $(\mathcal{C}, \mathcal{D})$-system. By Remark 8 we can find an infinite set I of indices such that for every $i \in I$: t_i has effective stack-height one and $a_i \in \Sigma_D$. For this we observe that if all configurations t_i with i greater than some number i_0 have effective stack-height one, we can take all indices $i \geq i_0$ such that $a_i = \top$; otherwise we take the set of indices such that t_i has effective stack-height one, but t_{i-1} does not – then $a_i \in \Sigma_D$. Since P_{fin} is finite and $|M_i| = n$ for all i, the set $\{M_i \mid i \in \mathbb{N}\}$ is finite. By the pigeonhole principle, there exist $i, j \in I$ such that $M_i = M_j$, $g_i = g_j$ and t_i, t_j have effective stack-height one, the same state and the same top stack symbol. In addition, we ask that $\top$ is performed in the run from (M_i, t_i, g_i) to (M_j, t_j, g_j).

We can then define a run of the desired form by repeating the part between i and j. We let $u = a_0 \cdots a_i$, $v = a_{i+1} \cdots a_j$, $M = M_i = M_j$, $g = g_i = g_j$. To define the configurations t_k, we observe that the configurations t_j and t_i can be represented as $t_i = qA\alpha$ and $t_j = qA\beta\alpha$. We can then define $t'_k = qA\beta^k\alpha$. We obtain a run of the $(\mathcal{C}, \mathcal{D})$-system: $(n[p_{init}^{fin}], t_{init}, g_{init}) \xrightarrow{u} (M, t'_1, g) \xrightarrow{v} (M, t'_2, g) \xrightarrow{v} \ldots$ □

4.2 Capacities and supported loops

The goal is a PSPACE algorithm for the existence of ultimately periodic runs in $(\mathcal{C}_{fin}, \mathcal{D})$. Since the reachability problem is decidable in PSPACE, we focus on loops. We follow the approach proposed in [21] for the reachability problem. Adapting this approach to infinite runs is not straightforward, and requires the new notion of ω-support.

As in [21], we decompose a $(\mathcal{C}_{fin}, \mathcal{D})$-system into a finite-state system $\mathcal{C}_{fin}^{\kappa}$ representing the contribution of $\mathcal{C}_{fin}$, and a pushdown system $\mathcal{D}^{\kappa}$ representing the contribution of $\mathcal{D}$.

The idea underlying the decomposition is the following. Once a value g has been written by a contributor into the register, replicating the contributor's run supplies arbitrary (but finitely) many contributor writes of g. This is captured by introducing a new set of actions $\Sigma_\nu = \{\nu(g) : g \in G\}$ denoting first contributor writes. In addition, each of $\mathcal{C}_{fin}^{\kappa}$ and $\mathcal{D}^{\kappa}$ have a component K called *capacity*, that stores the "writing" capacity that contributors already provided. Formally, the set of control states of $\mathcal{D}^{\kappa}$ is $\mathcal{P}(G) \times Q \times G$, and the initial state is $(\emptyset, q_{init}, g_{init})$. The input and the stack alphabets, Σ_D and Γ_D, are inherited from $\mathcal{D}$. So a configuration of $\mathcal{D}^{\kappa}$ has the form

$(K \subseteq G,\ t \in Q\Gamma_D^*,\ g \in G)$. The transitions of $\mathcal{D}^\kappa$ are:

$$
\begin{aligned}
&(K,t,g) \xrightarrow{w(h)} (K,t',h) && \text{if } t \xrightarrow{w(h)} t' \text{ in } \Delta,\\
&(K,t,g) \xrightarrow{r(h)} (K,t',h) && \text{if } t \xrightarrow{r(h)} t' \text{ in } \Delta \text{ and } h \in K \cup \{g\},\\
&(K,t,g) \xrightarrow{\nu(h)} (K \cup \{h\},t,h) && \text{if } h \notin K\,.
\end{aligned}
$$

The finite transition system $\mathcal{C}^\kappa_{fin}$ is defined similarly, it just follows in addition the transitions of $\mathcal{D}^\kappa$. The set of states of $\mathcal{C}^\kappa_{fin}$ is $\mathcal{P}(G) \times P_{fin} \times G$, input alphabet Σ_C, and initial state $(\emptyset, p^{fin}_{init}, g_{init})$. The transition relation δ^κ is:

$$
\begin{aligned}
&(K,p,g) \xrightarrow{w(h)} (K,p,h)\\
&(K,p,g) \xrightarrow{r(h)} (K,p,h)\\
&(K,p,g) \xrightarrow{\nu(h)} (K \cup \{h\},p,h)\\
&(K,p,g) \xrightarrow{\overline{w}(h)} (K,p',h) \quad \text{if } p \xrightarrow{\overline{w}(h)} p' \text{ in } \delta \text{ and } h \in K\\
&(K,p,g) \xrightarrow{\overline{r}(h)} (K,p',h) \quad \text{if } p \xrightarrow{\overline{r}(h)} p' \text{ in } \delta \text{ and } h \in K \cup \{g\}\,.
\end{aligned}
$$

Note that in both $\mathcal{D}^\kappa$ and $\mathcal{C}^\kappa$ some additional reads $r(h), \overline{r}(h)$ are possible when $h \in K$ – these are called *capacity reads*.
Notation. We write $\Sigma_{D,\nu}$ for $\Sigma_D \cup \Sigma_\nu$. Similarly for $\Sigma_{C,\nu}$ and $\Sigma_{C,D,\nu}$. By $v|_\Sigma$ we will denote the subword of v obtained by erasing the symbols not in Σ. Note that the value of the register after executing a trace v, in both $\mathcal{C}^\kappa_{fin}$ and $\mathcal{D}^\kappa$, is determined by the last action of v. We denote by $last(v)$ the register value of the last action of v (for v non-empty).

We now come back to examining when there exists an ultimately periodic run of the $(\mathcal{C}_{fin}, \mathcal{D})$-system, and focus on loops. Clearly, a loop in the $(\mathcal{C}_{fin}, \mathcal{D})$-systems leads to a loop in $\mathcal{D}^\kappa$, but the converse is not true. To recover the equivalence, we introduce the notion of ω-supported traces. Informally, a loop v of $\mathcal{D}^\kappa$ will be called supported when (1) for each $\nu(h)$ move in v there is a trace of $\mathcal{C}^\kappa_{fin}$ witnessing the fact that a contributor run can produce the required action $\overline{w}(h)$, and (2) all the witness traces can be completed to loops in $\mathcal{C}^\kappa_{fin}$.

Definition 10 Consider a word

$$v = v_1\nu(h_1)\cdots v_m\nu(h_m)v_{m+1} \in \Sigma^*_{D,\nu},$$

where $v_1, \ldots, v_{m+1} \in \Sigma^*_D$, and $h_1, \ldots, h_m$ are pairwise different register values. We say that v is *ω-supported* from $p_1, \ldots, p_m \in P_{fin}$ if for every $1 \leq i \leq m$ there is a word $u^i \in (\Sigma_{C,D,\nu})^*$ of the form

$$u^i = u^i_1\nu(h_1)\cdots u^i_i\nu(h_i)\boldsymbol{\overline{w}(h_i)}u^i_{i+1}\cdots u^i_m\nu(h_m)u^i_{m+1}$$

such that: (i) $u^i|_{\Sigma_{D,\nu}} = v$, and (ii) $(\emptyset, p_i, g) \xrightarrow{u^i} (K, p_i, g)$ in $\mathcal{C}^\kappa_{fin}$, where $g = last(v)$.

Note that $K = \{h_1, \ldots, h_m\}$ in the above definition, and that $u^i_j|_{\Sigma_{D,\nu}} = v_j$ holds for all j.

Lemma 11 The $(\mathcal{C}, \mathcal{D})$-system has a Büchi run iff there is some reachable configuration $(M, qA\alpha, g)$ in the $(\mathcal{C}_{fin}, \mathcal{D})$-system and a word $v \in \Sigma^*_{D,\nu}$ such that:

1. $\mathcal{D}^\kappa$ has a run of the form $(\emptyset, qA, g) \xrightarrow{v} (K, qA\alpha', g)$, and $\top$ appears in v.
2. v is ω-supported from some $p_1, \ldots, p_m$ such that $[p_1, \ldots, p_m] \leq M$.

Observe that by Definition 10, we have $m \leq |G|$ in Lemma 11.

Proof. We outline the proof, focussing on the right to left direction. We construct an infinite run of the $(\mathcal{C}_{fin}, \mathcal{D})$-system starting in $(M, qA\alpha, g)$ by shuffling in a suitable way (infinitely many copies of) the run v of $\mathcal{D}^\kappa$ and a number of copies of the runs of $\mathcal{C}^\kappa_{fin}$ supporting v.

Consider one of the $\nu(h)$ occuring in v. Since v is ω-supported, there is a run $\rho : p \xrightarrow{u_1} p_1 \xrightarrow{\overline{w}(h)} p_2 \xrightarrow{u_2} p$ of $\mathcal{C}^\kappa_{fin}$ that can be executed with v to produce the first occurrence of $\overline{w}(h)$, in place of $\nu(h)$. To execute v once, we can replicate the initial part $p \xrightarrow{u_1} p_1$ as many times as needed, and simulate each capacity read $r(h)$ or $\overline{r}(h)$ occuring later, in v or one of its supporting runs, by letting a different copy take the transition $p_1 \xrightarrow{\overline{w}(h)} p_2$ to enable the read. At the end of the first simulation of v, the main contributor executing ρ is back in state p, and the others are still in state p_2. During the simulation of the second loop, we use a second set of copies of the main contributor to produce in the same way all required write operations $\overline{w}(h)$. Meanwhile, the first set of contributors can be brought back to state p: as soon as the main contributor reaches again state p_2, the first set of copies resumes the run $p_2 \xrightarrow{u_2} p$ by following the run of the main contributor. □

4.3 Final step

As in the case of reachability, we show that $\mathcal{D}^\kappa$ can be replaced by a finite-state system representing its downward closure, since adding some transitions of the leader does not affect the support of contributors. This finite-state system will be synchronized with the contributor automata witnessing support, yielding the PSPACE algorithm.

Lemma 12 Let $v = v_1\nu(h_1)\cdots v_{m+1}$be ω-supported from $p_1, \ldots, p_m$, and let $v_j \sqsubseteq \overline{v}_j$ for every j. Assume that $\overline{v} = \overline{v}_1\nu(h_1)\cdots\overline{v}_{m+1}$ satisfies $last(v) = last(\overline{v})$. Then $\overline{v}$ is also ω-supported from $p_1, \ldots, p_m$.

Proof. We need to lift all traces u^i of $\mathcal{C}^\kappa_{fin}$ that witness that v is ω-supported, to traces $\overline{u}^i$ that ω-support $\overline{v}$. We assume that the traces u^i satisfy all assumptions of Definition 10, and write $g = last(v) = last(\overline{v}) = last(u^i)$.

Let $v_i = a_{i,1} \cdots a_{i,n_i}$ and $\overline{v}_i = x_{i,0}a_{i,1}x_{i,1} \cdots a_{i,n_i}x_{i,n_i}$, for some $x_{i,j} \in \Sigma_D^*$. We obtain $\overline{u}^i$ from u^i by substituting each $a_{i,j}$ by $x_{i,j-1}a_{i,j}$, each $\nu(h_l)$ by $x_{l,n_l}\nu(h_l)$ and adding $x_{m+1,n_{m+1}}$ at the end. By construction we have $\overline{u}^i|_{\Sigma_{D,\nu}} = \overline{v}$, and $last(\overline{u}^i) = g$, since $\overline{u}^i$ ends either with the same action as $x_{m+1,n_{m+1}}$ (i.e., as $\overline{v}$), or as u^i. Observe also that $(\emptyset, p_i, g) \xrightarrow{\overline{u}^i} (K, p_i, g)$ still holds in $\mathcal{C}^\kappa_{fin}$: actions of $\mathcal{D}$ can only modify the register component in $\mathcal{C}^\kappa_{fin}$, and this is the same after reading $a_{i,j}$ or $x_{i,j-1}a_{i,j}$. □

Combining known results for the reachability problem in $(\mathcal{C}, \mathcal{D})$-systems with Lemmas 11 and 12, we obtain a polynomial space algorithm for the repeated reachability problem:

Theorem 13 *The repeated reachability problem for $(\mathcal{C}, \mathcal{D})$-systems is* PSPACE*-complete when $\mathcal{C}$ and $\mathcal{D}$ range over pushdown systems.*

Proof. The lower bound follows from [12].

For the upper bound we introduce one more shorthand. Let $h_1, \ldots, h_m$ be a sequence of values from G. A $(h_1, \ldots, h_m)$-word is a word of the form $v_1\nu(h_1) \cdots v_m\nu(h_m)v_{m+1}$ with no occurrence of ν in $v_1 \cdots v_{m+1}$.

Our PSPACE algorithm consists of three steps:

1. Guess a sequence $h_1, \ldots, h_m$ of pairwise distinct values from G, states $p_1, \ldots, p_m$ of $\mathcal{C}_{fin}$, control state $q \in Q$ and stack symbol $A \in \Gamma_D$ for $\mathcal{D}$, and a value $g \in G$.

2. Check that there exists a configuration $(M, qA\alpha, g)$ satisfying $M \geq [p_1, \ldots, p_m]$ that is reachable in the $(\mathcal{C}_{fin}, \mathcal{D})$-system.

3. Check that there exists a $(h_1, \ldots, h_m)$-word $v \in \Sigma^*_{D,\nu}$ with $last(v) = g$ such that:

   (a) v is ω-supported from $p_1, \ldots, p_m$,

   (b) $\mathcal{D}^\kappa$ has a run of the form $(\emptyset, qA, g) \xrightarrow{\overline{v}} (K, qA\alpha', g)$ for some $(h_1, \ldots, h_m)$-word $\overline{v}$ such that $v \sqsubseteq \overline{v}$, $last(\overline{v}) = last(v) = g$, and $\top$ occurs in $\overline{v}$.

By Lemmas 11 and 12, this algorithm returns "yes" iff the system has a Büchi run.

Step 1 can be done in polynomial space since we have $m \leq |G|$.

Step 2 reduces to an instance of the reachability problem in the $(\mathcal{C}, \mathcal{D})$-system. First, by Lemma 7, the question is to decide if there exists a configuration $(M, qA\alpha, g)$ in the $(\mathcal{C}, \mathcal{D})$-system that is reachable and such that

$M \geq [p_1\alpha_1, \ldots, p_m\alpha_n]$ for some $\alpha_1, \ldots, \alpha_n \in \Gamma^*$. We modify $\mathcal{C}$ and $\mathcal{D}$ so that after such a configuration has been reached (and only in that case), the leader can do a new action $\top_0$. The idea is to add transitions to $\mathcal{C}$ so that a contributor in a state of the form $p_i\alpha$ can write "i" to the register (once), and add transitions $qA \xrightarrow{r(1)} \ldots \xrightarrow{r(m)} \xrightarrow{\top_0}$ to $\mathcal{D}$. Since the reachability problem is in PSPACE when $\mathcal{D}$ and $\mathcal{C}$ are pushdown systems [12, 21], step 2 is in PSPACE.

Step 3 requires to construct some auxiliary automata. For $i = 1, \ldots, m$, let $\mathcal{A}_i$ be a finite automaton accepting the projection over $\Sigma_{D,\nu}$ of the words $u \in \Sigma^*_{D,C,\nu}$ of the form

$$u = u_1\nu(h_1)\cdots u_i\nu(h_i)\overline{\boldsymbol{w}}(\boldsymbol{h_i})u_{i+1}\cdots u_m\nu(h_m)u_{m+1} \tag{3}$$

and such that $(\emptyset, p_i, g) \xrightarrow{u} (K, p_i, g)$ is a trace in $\mathcal{C}^\kappa_{fin}$. Since $\mathcal{C}^\kappa_{fin}$ can be computed in PSPACE, so does $\mathcal{A}_i$.

Consider the language:

$$L = \{\overline{v} \in \Sigma^*_{\mathcal{D},\nu} \ : \ \overline{v} \text{ contains } \top,\ last(\overline{v}) = g,\ \overline{v} \text{ is a } (h_1, \ldots, h_m)\text{-word}, \\ (\emptyset, qA, g) \xrightarrow{\overline{v}} (K, qA\alpha, g) \text{ in } \mathcal{D}^\kappa, \text{ for some } \alpha\}\,.$$

A pushdown automaton recognizing L can be obtained by slightly modifying $\mathcal{D}^\kappa$. Since L can be recognized by a pushdown automaton of polynomial size, it is possible to compute on-the-fly in PSPACE a finite automaton $\mathcal{A}$ (of exponential size) accepting all the $(h_1, \ldots, h_m)$-words in the downward closure of L, see [8].

Step 3 then consists in checking that the intersection of the automata $\mathcal{A}$, $\mathcal{A}_1, \ldots, \mathcal{A}_m$ is non-empty, which can be done in PSPACE. □

Remark 14 In the case where there is no leader, $\mathcal{D}^\kappa$ becomes an automaton accepting all sequences $\nu(h_1)\cdots\nu(h_m)$ such that $h_1, \ldots, h_m$ are pairwise distinct. To check that such a sequence is ω-supported, we can test separately for each $\nu(h_i)$ if there is a contributor run that can produce $\overline{w}(h_i)$ – instead of having to take the product of m automata in order to synchronize with $\mathcal{D}^\kappa$, as in the proof of Theorem 13. For that reason, we can test directly for each $\nu(h_i)$ if the pushdown automaton $\mathcal{C}^\kappa$ (rather than $\mathcal{C}^\kappa_{fin}$) admits a run as in Equation 3. This leads to an algorithm in NP instead of PSPACE.

5 Max-safe problem

We show in this section that the max-safe problem is NP-complete when $\mathcal{C}$ ranges over finite-state systems and $\mathcal{D}$ ranges over pushdown systems, and NEXPTIME-complete when both $\mathcal{C}$ and $\mathcal{D}$ range over pushdown systems (Theorem 4).

We start by introducing a set semantics of $(\mathcal{C}, \mathcal{D})$-systems, that replaces multisets by sets. This semantics is suitable for the reachability and max-safe problems (but not for liveness). We show that the max-safe problem is NP-complete when contributors are finite-state. Then we consider the case of contributors given by a pushdown automaton. As for liveness we can reduce this case to the case when contributors are finite-state. This gives a NEXPTIME algorithm.

5.1 Set semantics

As a first step we will introduce the set semantics of $(\mathcal{C}, \mathcal{D})$-systems that is equivalent to the multiset semantics of Section 2 when only finite traces are considered. The idea is that since the number of contributors is arbitrary, we can always add contributors that copy all the actions of a given contributor. So once a state of $\mathcal{C}$ is reached, we can assume that we have arbitrarily many copies of $\mathcal{C}$ in that state. In consequence, we can replace multisets by sets. A very similar semantics has already been used in [21, 10]. Here we need to be a bit finer in order to handle deadlocks.

Consider a $(\mathcal{C}, \mathcal{D})$-system with the notations as in Eq. (1) on page 4:

$$\mathcal{C} = \langle S, \delta, s_{init} \rangle \qquad \mathcal{D} = \langle T, \Delta, t_{init} \rangle \ .$$

Instead of multisets $M \in \mathbb{N}^S$, we use sets $B \subseteq S$. As for multisets we lift the transitions from elements to sets of elements:

$$B \xrightarrow{a} B' \text{ in } \delta \quad \text{if } s \xrightarrow{a} s' \text{ in } \delta, \text{ and } B' \text{ is either } B \cup \{s'\} \text{ or } (B \cup \{s'\}) \setminus \{s\} \text{ for some } s \in B.$$

The intuition is that $B \xrightarrow{a} B \cup \{s'\}$ represents the case where *some* contributors in state s take the transition, and $B \xrightarrow{a} (B \cup \{s'\}) \setminus \{s\}$ corresponds to the case where *all* contributors in state s take the transition. The transitions in the *set semantics* are essentially the same as for the multiset case:

$$\begin{aligned}
(B, t, g) &\xrightarrow{w(h)} (B, t', h) && \text{if } t \xrightarrow{w(h)} t' \text{ in } \Delta \\
(B, t, g) &\xrightarrow{r(h)} (B, t', h) && \text{if } t \xrightarrow{r(h)} t' \text{ in } \Delta \text{ and } h = g \\
(B, t, g) &\xrightarrow{\overline{w}(h)} (B', t, h) && \text{if } B \xrightarrow{\overline{w}(h)} B' \text{ in } \delta \\
(B, t, g) &\xrightarrow{\overline{r}(h)} (B', t, h) && \text{if } B \xrightarrow{\overline{r}(h)} B' \text{ in } \delta \text{ and } h = g
\end{aligned}$$

Lemma 15 1. If $(M_0, t_0, g_0) \xrightarrow{a_1} \ldots \xrightarrow{a_n} (M_n, t_n, g_n)$ in the multiset semantics, and B_j is the support of M_j, for every $j = 0, \ldots, n$, then $(B_0, t_0, g_0) \xrightarrow{a_1} \ldots \xrightarrow{a_n} (B_n, t_n, g_n)$ in the set semantics.

2. If $(B_0, t_0, g_0) \xrightarrow{a_1} \ldots \xrightarrow{a_n} (B_n, t_n, g_n)$ in the set semantics, then there exist multisets $M_0, \ldots, M_n$ such that M_j has support B_j, and for some $i_j > 0$,

$$(M_0, t_0, g_0) \xrightarrow{(a_1)^{i_1}} (M_1, t_1, g_1) \xrightarrow{(a_2)^{i_2}} \ldots \xrightarrow{(a_n)^{i_n}} (M_n, t_n, g_n)$$

in the multiset semantics.

Proof. The first part of the lemma follows directly from the definitions. For the second part, writing $B_0 = \{s_1, \ldots, s_n\}$, we let $M_0 = 2^n[s_1, \ldots, s_m]$. We then simulate a step of the set semantics by letting either the leader, all the copies of $\mathcal{C}$, or half the copies of $\mathcal{C}$ take the transition in the $(\mathcal{C}, \mathcal{D})$-system. □

Remark 16 The set semantics is a variant of the *accumulator semantics* used in [21], in which only transitions of the form $B \xrightarrow{a} B \cup \{s'\}$ (but not $B \xrightarrow{a} (B \cup \{s'\}) \setminus \{s\}$) were used. The accumulator semantics is sufficient for the reachability problem, and has the nice property that the B-part is monotonic (hence the name *accumulator* semantics). So for instance, in the case of finite-state processes, it leads to a very simple NP algorithm for the reachability problem [21]. However, the accumulator semantics does not satisfy the first item of Lemma 15, and is thus not precise enough for properties that refer to *maximal* runs.

Corollary 17 Fix a $(\mathcal{C}, \mathcal{D})$-system. In the multiset semantics the system has a finite maximal safe run ending in a configuration (M, t, g) iff in the set semantics the system has a finite maximal safe run ending in the configuration (B, t, g) with B being the support of M.

5.2 Finite-state contributors

First we will assume that $\mathcal{C}$ is a finite-state transition system.

Lemma 18 The max-safe problem is NP-hard when $\mathcal{C}$ and $\mathcal{D}$ are both finite-state.

Proof. We reduce 3-SAT to the max-safe problem. Given a formula $\varphi = c_1 \wedge \ldots \wedge c_m$, with c_j clauses of length 3 over variables $x_1, \ldots, x_n$, we construct finite-state processes $\mathcal{C}$ and $\mathcal{D}$ such that φ is satisfiable iff there exists a maximal run in the $(\mathcal{C}, \mathcal{D})$-system that contains no occurrence of a fixed action $\top$ of $\mathcal{D}$.

The leader $\mathcal{D}$ will guess the values of the variables, and the contributors will check if all clauses are satisfied. So, the leader starts by successively writing $x_1 = b_1, \ldots, x_n = b_n$ in the register, where each b_i is a guessed truth value. Meanwhile, each contributor chooses a clause c_j, reads the values of

the variables appearing in c_j as the leader writes them, and checks whether c_j is satisfied. If it is, he writes "c_j" in the register. After having guessed values for the all the variables, the leader must successively read "c_1", ..., "c_m". If she manages to do this then she knows that all clauses are satisfied. She enters then some state q with no outgoing transitions. For every other state $q' \neq q$ of the leader we add a transition $q' \xrightarrow{\top} q'$.

Suppose φ is satisfiable. Take a run as described above, starting with n contributors, and where the leader chooses a valuation that satisfies φ. The leader ends in state q, from which she has no available transition. Similarly, each contributor stops after writing one of the "c_i", or blocks because he missed reading some variable. So the run is maximal, and does not contain $\top$.

For the other direction, observe that all safe runs must be finite and should end with $\mathcal{D}$ in the state q because all other states have a transition on $\top$. Such a run defines a valuation of the variables that satisfies φ. □

To decide the max-safe problem, we check separately for the existence of an infinite, or finite and maximal, run without occurrences of $\top$ (a *safe run*). The case of infinite safe runs can be reduced to the repeated reachability problem: by Theorem 5 we can construct from $(\mathcal{C}, \mathcal{D})$ an equivalent $(\tilde{\mathcal{C}}, \tilde{\mathcal{D}})$-system in which all infinite runs contains infinitely many writes from the leader. To decide if this system admits an infinite run, we can then test for each possible value g of the register if there is a run with infinitely many occurrences of $w(g)$. Since the repeated reachability problem is in NP for finite-state contributors [10] we obtain:

Lemma 19 When $\mathcal{C}$ ranges over finite-state systems and $\mathcal{D}$ over pushdown systems, deciding whether a $(\mathcal{C}, \mathcal{D})$-system has an infinite safe run is in NP.

It remains to give an algorithm for the existence of a *finite* maximal safe run. By Corollary 17 we can use the set semantics. From now on we will also need to exploit the fact that $\mathcal{D}$ is a pushdown system. Recall that the states of $\mathcal{D}$ are of the form $q\alpha$ where q is the state of the pushdown automaton defining $\mathcal{D}$ and α represents the stack. The question is to decide if there is a configuration $(B, q\alpha, g)$ from which there is no outgoing transition in the $(\mathcal{C}, \mathcal{D})$-system, and such that $(B, q\alpha, g)$ is reachable without using $\top$ actions. Note that we can say whether $(B, q\alpha, g)$ has no outgoing transition by looking only at B, q, g and the top symbol of α. Our algorithm will consists in guessing B, q, g and some $A \in \Gamma_D$, and checking reachability. First, we show that it is sufficient to look for traces where the number of changes to the first component of configurations is bounded. The idea is that we can always assume that in a run of the $(\mathcal{C}, \mathcal{D})$-system, a state s is added at most once to the current set of contributor states B. This simply means that a state is removed from B only if it will never be used again in the run.

Lemma 20 Let $\rho = (B_0, t_0, g_0) \xrightarrow{a_1} (B_1, t_1, g_1) \xrightarrow{a_2} \dots \xrightarrow{a_n} (B_n, t_n, g_n)$ be a run of the $(\mathcal{C}, \mathcal{D})$-system. There exists a run $\rho' = (B'_0, t_0, g_0) \xrightarrow{a_1} (B'_1, t_1, g_1) \xrightarrow{a_2} \dots \xrightarrow{a_n} (B'_n, t_n, g_n)$ such that $B_0 = B'_0$, $B_n = B'_n$, and for all $s \in S$ and $0 \leq i < n$, if $s \in B'_i$ and $s \notin B'_{i+1}$, then for all $j > i$, $s \notin B'_j$.

Proof. We define B'_i by induction on i: $B'_0 = B_0$, and for $i > 1$,

- if $B_{i+1} = B_i$, then $B'_{i+1} = B'_i$.
- if $B_{i+1} = B_i \cup \{s\}$, then $B'_{i+1} = B'_i \cup \{s\}$.
- if $B_{i+1} = (B_i \setminus \{s\}) \cup \{s'\}$ and $s \notin B_j$ for all $j > i$, then $B'_{i+1} = (B'_i \setminus \{s\}) \cup \{s'\}$. If $s \in B_j$ for some $j > i$, then $B'_{i+1} = B'_i \cup \{s'\}$.

Clearly, ρ' is a run of the $(\mathcal{C}, \mathcal{D})$-system. Moreover, for all i, $B_i \subseteq B'_i \subseteq \bigcup_{j=i}^{n} B_j$. So in particular, $B_n = B'_n$. □

Corollary 21 Every finite run ρ of the $(\mathcal{C}, \mathcal{D})$-system in the set semantics can be written as $\rho = \rho_0 \cdots \rho_k$ with $k \leq 2|S|$, where in each ρ_j, all states have the same first component.

Proof. We take a run of the form described in Lemma 20. Let $i_0 = 0$, and $i_1 < \cdots < i_k$ be the indices such that $B_i \neq B_{i-1}$. Consider the sequence $B_{i_0}, B_{i_1}, \dots, B_{i_k}$. There are states $s_1, \dots, s_k \in S$ such that for all $0 \leq j < k$, $B_{i_{j+1}} = B_{i_j} \cup \{s_j\}$ or $B_{i_{j+1}} = (B_{i_j} \cup \{s\}) \setminus \{s_j\}$ for some s. Moreover, each $s \in S$ is added at most once, and removed at most once from some B_i, which means that there are at most two distinct indices j such that $s = s_j$. Hence $k \leq 2|S|$. □

Lemma 22 For every finite run ρ of the $(\mathcal{C}, \mathcal{D})$-system in the set semantics, there exists a run ρ' with same label and end configuration that can be written as $\rho' = \rho_0 \cdots \rho_k$ with $k \leq 2|S|$, where in each ρ_j, all states have the same first component.

Proof. We take a run of the form described in Lemma 20. Let $i_0 = 0$, and $i_1 < \cdots < i_k$ be the indices such that $B_i \neq B_{i-1}$. Consider the sequence $B_{i_0}, B_{i_1}, \dots, B_{i_k}$. There are states $s_1, \dots, s_k \in S$ such that for all $0 \leq j < k$, $B_{i_{j+1}} = B_{i_j} \cup \{s_j\}$ or $B_{i_{j+1}} = (B_{i_j} \cup \{s\}) \setminus \{s_j\}$ for some s. Moreover, each $s \in S$ is added at most once, and removed at most once from some B_i, which means that there are at most two distinct indices j such that $s = s_j$. Hence $k \leq 2|S|$. □

Lemma 23 The following problem belongs to NP:
Input: finite-state system $\mathcal{C}$, pushdown system $\mathcal{D}$, $B \subseteq S$, $q \in Q$, $A \in \Gamma_D$.
Question: Does the $(\mathcal{C}, \mathcal{D})$-system admit a run from $(\{s_{init}\}, q_{init} A^D_{init}, g_{init})$ to $(B, qA\alpha, g)$ for some $\alpha \in \Gamma_D^*$?

Proof. The set semantics of the $(\mathcal{C}, \mathcal{D})$-system can be described by a pushdown automaton $\mathcal{A}$ with set of control states $2^S \times Q \times G$, input alphabet $\Sigma_C \cup \Sigma_D$, and stack alphabet Γ_D. We guess a sequence $\{s_{init}\} = B_0, B_1, \ldots, B_k = B$ where $k \leq 2|S|$, and construct the restriction of the pushdown automaton $\mathcal{A}$ to runs where the first component of the state takes values according to $B_0, B_1, \ldots, B_k$. This new pushdown automaton is of polynomial size, and we can check whether it has a reachable configuration $(B, qA\alpha, g)$ in polynomial time [4]. □

By Lemmas 19 and 23, the max-safe problem is thus in NP and we obtain:

Theorem 24 *The max-safe problem is* NP*-complete when $\mathcal{C}$ ranges over finite-state systems and $\mathcal{D}$ ranges over pushdown systems.*

5.3 Pushdown contributors

We now return to the case where both $\mathcal{C}$ and $\mathcal{D}$ are given by pushdown systems.

Lemma 25 The max-safe problem is NEXPTIME-hard when $\mathcal{C}$ and $\mathcal{D}$ range over pushdown systems.

Proof. We reduce the following tiling problem to the max-safe problem:

Input: A finite set of tiles Σ, horizontal and vertical compatibility relations $H, V \subseteq \Sigma^2$, and initial row $x \in \Sigma^n$.
Question: is there a tiling of the $2^n \times 2^n$ square respecting the compatibility relations and containing the initial row in the left corner?

A tiling is a function $t : \{1, \ldots, 2^n\}^2 \to \Sigma$ such that $(t(i,j), t(i,j{+}1)) \in H$ and $(t(i,j), t(i+1,j)) \in V$ for all i, j, and $t(1,1)t(1,2)\cdots t(1,n) = x$.

The idea of the reduction is that the system will have a maximal run without $\top$ if and only if the leader guesses a tiling respecting the horizontal compatibility, and the contributors check that the vertical compatibility is respected as well.

The leader will write down the tiling from left to right and from top to bottom, starting with the initial row. The sequence of values taken by the register on a (good) run will have the form

$$A_{1,1}, \overline{A_{1,1}}, A_{1,2}, \overline{A_{1,2}}, \ldots, A_{1,2^n}, \overline{A_{1,2^n}}, \ldots, A_{2^n,2^n}\ \overline{A_{2^n,2^n}}\ (\$\overline{\$})^{2^n} \diamond\ .$$

The $A_{i,j}$ are guessed and written by the leader, and the $\overline{A_{i,j}}$ are written by contributors. Letters $\overline{A_{i,j}}$ have two purposes: they ensure that at least one contributor has read the preceding letter, and prevent a contributor to read the same letter twice. For technical reasons, this sequence is followed by a sequence $(\$\overline{\$})^{2^n}\diamond$ of writes from the leader (with $\$, \diamond \notin \Sigma$), and we will consider that $(A, \$) \in V$ for all $A \in \Sigma$.

The leader uses her stack to count the number i of rows (using the lower part of the stack), and the number j of tiles on each row (using the upper part of the stack). So, she repeats the following, up to reaching the values $i = 2^n, j = 2^n$:

- guess a tile A compatible with the one on its left (if $j \neq 1$), and write A on the register,
- wait for an acknowledgment $\overline{A}$ from one of the contributors,
- increment j,
- if $j > 2^n$, increment i and set $j = 1$.

Finally, she repeats 2^n times the actions $w(\$)$, $w(\overline{\$})$, then finishes by writing $w(\diamond)$ and going to state q_f.

Each contributor is supposed to read the entire sequence of values written in the register. He alternates between reading values of the form A and $\overline{A}$, which ensures that no value is read more than one time. At the same time, he uses his stack to count the number of writes $w(A)$ ($A \in \Sigma \cup \{\$\}$) of the leader, up to $(2^{2n} + 2^n)$, so that he can check that no value was missed. This operation will in fact be divided between counting up to 2^{2n}, and counting up to 2^n, as described below.

Every contributor decides non-deterministically to check vertical compatibility at some point. He chooses the current tile $A \neq \$$, and needs to check that the tile located below it (that is, occurring 2^n tiles later in the sequence of values written by the leader) is compatible with it. This is done as follows: after reading $A \neq \$$, the contributor writes $\overline{A}$ on the register (rather than waiting for another contributor to do so), and remembers the value. He interrupts his current counting, and starts counting anew on the top of the stack, up to 2^n. Upon reaching 2^n, he stores the value A' of the register, for later check. Then he resumes the first counting while reading the remaining of the sequence, up to 2^{2n}. At any moment, the contributor can read $\diamond$. If he reads $\diamond$ and either $(A, A') \notin V$ or the counting up to 2^n failed (i.e., his stack is not empty), then he writes $\# \notin G$ and stops; otherwise he simply stops. In state q_f, the leader may read any value $g \neq \diamond$, and she then do $\top$: $q_f \xrightarrow{r(g)} \xrightarrow{\top}$. From every other state $q \neq q_f$, the leader can do $\top$, too.

If there is a tiling of the $2^n \times 2^n$ square, then we obtain a maximal run with 2^{2n} contributors and without any occurrence of $\top$, by letting the leader write the sequence of register values corresponding to this tiling, and having each contributor perform one of the 2^{2n} vertical compatibility checks. If each contributor reads every value produced by the leader, his stack will be empty upon reading $\diamond$, so he simply stops and no $\top$ will be generated.

Conversely, we show that in any maximal run without $\top$, the sequence of tiles guessed by the leader defines a correct tiling of the $2^n \times 2^n$ square.

First, in any such run the leader needs to reach state q_f: if she gets no acknowledgment on some $A \in \Sigma$ then she would do $\top$, which is impossible by assumption. So she guesses a sequence $A_{1,1}, \ldots, A_{2^n,2^n}$ with $(A_{i,j}, A_{i,j+1}) \in H$ for all $j < 2^n$, gets an acknowledgment $\overline{A}_{i,j}$ for each $A_{i,j}$, and finally writes $(\$\overline{\$})^{2^n} \diamond$. Moreover, $\diamond$ is the final value of the register: if it were overwritten by a contributor, the leader could generate $\top$. So every contributor that has chosen some $A_{i,j} \in \Sigma$ will ultimately read $\diamond$. Since he cannot do $\overline{w}(\#)$, his stack must be empty at that point and he must have successfully checked that $(A_{i,j}, A_{i+1,j}) \in V$. □

As with finite-state contributors, to solve the max-safe problem, we look separately for an infinite safe run, or a finite maximal safe run. The case of infinite runs can again be reduced to the repeated reachability problem, using Theorem 5.

Lemma 26 When $\mathcal{C}$ and $\mathcal{D}$ range over pushdown systems, deciding whether a $(\mathcal{C}, \mathcal{D})$-system has an infinite safe run is in PSPACE.

To decide the existence of a *finite* maximal safe run, we reduce the problem to the case of finite-state contributors, using Lemma 7.

Lemma 27 When $\mathcal{C}$ and $\mathcal{D}$ range over pushdown systems, deciding whether a $(\mathcal{C}, \mathcal{D})$-system has a finite maximal safe run is in NEXPTIME.

Proof. We define the *top* of a configuration of the $(\mathcal{C}, \mathcal{D})$-system as follows:

$$top(\{p_1 A_1 \alpha_1, \ldots, p_n A_n \alpha_n\}, qA\alpha, g) = (\{p_1 A_1, \ldots, p_n A_n\}, qA, g)$$

Observe that to determine if a configuration in the set semantics is a deadlock or not it suffices to look at its top. Moreover, by Lemma 15, deadlocks occur in the multiset semantics iff they occur in the set semantics.

The algorithm to decide the existence of a maximal safe run is as follows: guess a configuration top that corresponds to deadlocks in the $(\mathcal{C}, \mathcal{D})$-system, and check if it is reachable after removing all $\top$-transitions from the $(\mathcal{C}, \mathcal{D})$-system. By Lemma 7, this amounts to deciding if it is reachable in the $(\mathcal{C}_{fin}, \mathcal{D})$-system. Applying Lemma 23 to the $(\mathcal{C}_{fin}, \mathcal{D})$-system, which is of exponential size, this can be done in NEXPTIME. □

Theorem 28 *The max-safe problem is* NEXPTIME*-complete when $\mathcal{C}$ and $\mathcal{D}$ range over pushdown systems.*

5.4 Universal reachability

Recall that the *universal reachability* problem asks if all maximal runs of a given $(\mathcal{C}, \mathcal{D})$-system, so for every number of contributors, contain some occurrence of a special action $\top$. Correctness problems for parametrized

distributed algorithms can be rephrased as instances of universal reachability: we want to know whether for an arbitrary number of participants, and for every run of the algorithm, the outcome is correct. Correctness of the outcome is expressed here by the leader executing the action $\top$.

Remark 29 A natural variant of the universal reachability problem would be the following: is there some bound N such that for all $n \geq N$, all maximal runs with n contributors contain an occurrence of $\top$? A bit surprisingly, this formulation is equivalent to the universal reachability problem: if there were some maximal run with $n < N$ contributors without $\top$, then we could add arbitrary many contributors doing the same actions as one original contributor, thus obtaining a maximal run with N contributors and without $\top$, contradiction.

Since the max-safe problem is the complement of universal reachability, we obtain from Theorems 24 and 28:

Corollary 30 The universal reachability problem is coNP-complete for $(\mathcal{C}, \mathcal{D})$-systems where $\mathcal{C}$ is finite-state and $\mathcal{D}$ is a pushdown system. It is coNEXPTIME-complete when both $\mathcal{C}$ and $\mathcal{D}$ are pushdown systems.

6 Regular $\mathcal{C}$-expanding properties

In this section, we prove our general result stated in Theorem 2.

The proof of Theorem 2 is (again) divided into two cases: one for finite and the other for infinite traces. For finite maximal traces we use the results about the max-safe problem, and for infinite traces we give a reduction to repeated reachability.

Lemma 31 It is NEXPTIME-complete to decide whether a given pushdown $(\mathcal{C}, \mathcal{D})$-system has a *finite* maximal trace satisfying some $\mathcal{C}$-expanding property $\mathcal{P}$ given by a finite automaton or an LTL formula.

Proof. By Theorem 5 we can assume that we deal with a property $\mathcal{P}_D$ referring only to actions of $\mathcal{D}$. If $\mathcal{P}_D$ is given by an LTL formula, we start by constructing an equivalent finite automaton of exponential size. By taking the product of $\mathcal{D}$ with this automaton, we can assume that $\mathcal{D}$ has a distinguished set of final (control) states such that a finite run of the $(\mathcal{C}, \mathcal{D})$-system satisfies $\mathcal{P}_D$ iff $\mathcal{D}$ ends in a final state.

The result then follows using Lemma 7, together with Lemma 23. Recall that in order to decide if a finite run is maximal it is enough to look at the top of its last configuration. Lemma 7 then tell us that there exists a maximal finite run in the $(\mathcal{C}, \mathcal{D})$-system with $\mathcal{D}$ ending in a final state iff there exists such a run in the $(\mathcal{C}_{fin}, \mathcal{D})$-system; and by Lemma 23 this can be decided in NP in the size of $(\mathcal{C}_{fin}, \mathcal{D})$, so overall in NEXPTIME. The matching

NEXPTIME-hardness lower bound follows from the proof of Lemma 25, as the $(\mathcal{C}, \mathcal{D})$-system constructed there has no infinite safe trace, and the max-safe problem restricted to finite traces is a special case of our problem. □

The case of infinite runs turns out to be easier complexity-wise: PSPACE if the property is given by an automaton, and EXPTIME if it is given by an LTL formula.

Lemma 32 It is PSPACE-complete to decide whether a given pushdown $(\mathcal{C}, \mathcal{D})$-system has an *infinite* maximal trace satisfying a $\mathcal{C}$-expanding property $\mathcal{P}$ given by a Büchi automaton.

Proof. Applying again Theorem 5 and slightly modifying the $(\mathcal{C}, \mathcal{D})$-system we can reduce the satisfaction of $\mathcal{P}$ to an instance of the repeated reachability problem. Observe also that the repeated reachability problem is a special case of our problem. With this reduction, PSPACE-completeness follows from Theorem 3.

Let the pushdown system for $\mathcal{D}$ be $\mathcal{A}_D = \langle Q, \Sigma_D, \Gamma_D, \Delta, q_{init}, A^D_{init} \rangle$. By taking the product of $\mathcal{D}$ with a Büchi automaton for $\mathcal{P}_D$, we can also assume that $\mathcal{D}$ has a distinguished set R of repeating (control) states such that an infinite run of the $(\mathcal{C}, \mathcal{D})$-system satisfies $\mathcal{P}$ iff $\mathcal{D}$ visits R infinitely often.

We add new states and transitions to $\mathcal{A}_D$, so that the leader will signal visits to R by writing a special symbol $\# \notin G$. We set $G' = G \cup \{\#\}$, and $Q' = Q \cup \hat{Q}$, where $\hat{Q} = \{\hat{q} \mid q \in Q\}$ is a copy of Q. The stack alphabet is unchanged, and we add the following transitions to $\mathcal{A}_D$:

1. $r \xrightarrow{a} \hat{q}$ for every $r \xrightarrow{a} q$ with $r \in R$,
2. $\hat{q}_1 \xrightarrow{r(g)} \hat{q}_2$ for every $q_1 \xrightarrow{r(g)} q_2$,
3. $\hat{q}_1 \xrightarrow{w(\#)w(g)} q_2$ for every $q_1 \xrightarrow{w(g)} q_2$.

Note that $w(\#)$ does not restrict runs of the original system, and does not add new behaviours: the value # cannot be read by contributors, and it is immediately followed by original writes of the leader. Since on every infinite run the leader does infinitely often writes, she will write # infinitely often iff she visits infinitely often a state from R. □

Lemma 33 It is EXPTIME-complete to decide whether a given pushdown $(\mathcal{C}, \mathcal{D})$-system has an *infinite* trace satisfying some $\mathcal{C}$-expanding property $\mathcal{P}$, that is given by an LTL formula.

Proof. The lower bound comes from the situation where there are no contributors at all [4].

For the upper bound: from an LTL formula we first construct a Büchi automaton of exponential size for $\mathcal{P}$. As in Lemma 32, the first step is to

reduce the problem of deciding if the $(\mathcal{C}, \mathcal{D})$-system has a trace in $\mathcal{P}$ to a repeated reachability problem in some $(\mathcal{C}', \mathcal{D}')$-system. The leader $\mathcal{D}'$ there is of exponential size, and $\mathcal{C}'$ is of polynomial size.

As a second step we adapt the procedure given in the proof of Theorem 13: we do not build the downward closure of the leader, but enumerate all possible sequences $\nu(h_1), \ldots, \nu(h_m)$ and intermediate states, instead of guessing them. Then we follow the lines of the proof of Theorem 13, checking emptiness of pushdowns of exponential size (in EXPTIME).

First, there are exponentially many possible values for tuples of the form $(h_1, \ldots, h_m, p_1, \ldots, p_m, q, A, g)$ where $m \leq |G|$, $h_1, \ldots, h_m$ is a sequence of pairwise distinct values from G, $p_1, \ldots, p_m$ are states of $(\mathcal{C}'_{fin})^\kappa$, q is a control state of $(\mathcal{D}')^\kappa$, A is a stack symbol of $(\mathcal{D}')^\kappa$, and $g \in G$.

Then, for each such tuple, we can check in (deterministic) exponential time if there exists a $(h_1, \ldots, h_m)$-word $v \in \Sigma^*_{\mathcal{D},\nu}$ with $last(v) = g$ such that $(\emptyset, qA, g) \xrightarrow{v} (K, qA\alpha', g)$ in $(\mathcal{D}')^\kappa$, for some α', and v is ω-supported from $p_1, \ldots, p_m$. As in Theorem 13, we construct for every $1 \leq i \leq m$ a finite automaton $\mathcal{A}_i$ accepting the projection over $\Sigma_{D,\nu}$ of the words $u \in \Sigma^*_{C,D,\nu}$ of the form

$$u = u_1\nu(h_1)\cdots u_i\nu(h_i)\overline{\boldsymbol{w}}(\boldsymbol{h_i})u_{i+1}\cdots u_m\nu(h_m)u_{m+1}$$

and such that $(\emptyset, p_i, g) \xrightarrow{u} (K, p_i, g)$ is a trace in $(\mathcal{C}'_{fin})^\kappa$. Let $\mathcal{A}$ be a pushdown automaton accepting the set of $(h_1, \ldots, h_m)$-words v such that $(\emptyset, qA, g) \xrightarrow{v} (K, qA\alpha, g)$ in $(\mathcal{D}')^\kappa$. To decide the existence of an ω-supported trace satisfying the above conditions, we construct $\mathcal{A} \cap \mathcal{A}_1 \cap \cdots \cap \mathcal{A}_m$, which is a pushdown automaton of exponential size, and test whether its language is empty.

The PSPACE procedure for the reachability problem described in [21] is very similar, and we can adapt it in the same way to decide in exponential time if there exists a configuration $(M, qA\alpha, g)$ satisfying $M \geq [p_1, \ldots, p_m]$ that is reachable in the $(\mathcal{C}'_{fin}, \mathcal{D}')$-system.

By Lemma 11, this gives us an EXPTIME algorithm to decide if the original $(\mathcal{C}, \mathcal{D})$-system has a trace in $\mathcal{P}$. □

7 Simplifying $(\mathcal{C}, \mathcal{D})$-systems

In this section, we show that a $(\mathcal{C}, \mathcal{D})$-system can be simulated by another $(\mathcal{C}', \mathcal{D}')$-system such that all actions in the original system are reflected in leader writes in the new system. So in the $(\mathcal{C}', \mathcal{D}')$-system all behaviours of the system, up to stuttering, will be reflected in the actions of the leader.

The idea is that in the $(\mathcal{C}', \mathcal{D}')$-system the register of the $(\mathcal{C}, \mathcal{D})$-system becomes part of the leader state. This releases the actual register of the $(\mathcal{C}', \mathcal{D}')$-system to be used to communicate about contributor actions. Contributors will write into the register the command they want to perform, and

the leader will execute the command and confirm it by writing back into the register. This confirmation is read by contributors who at this point know that their request has been read and executed. For symmetry the leader is also writing the commands she performs to the register (although they are never read by anybody). So the set of *register values* of the $(\mathcal{C}', \mathcal{D}')$-systems is:

$$G' = \{\overline{\mathtt{r?}}(g), \overline{\mathtt{w?}}(g), \overline{\mathtt{r}}(g), \overline{\mathtt{w}}(g), \mathtt{r}(g), \mathtt{w}(g) : g \in G\} \cup \{g'_{init}\} \tag{4}$$

The alphabets of $\mathcal{C}'$ and $\mathcal{D}'$ are defined as usual:

$$\Sigma'_C = \{\overline{r}(g'), \overline{w}(g') : g' \in G'\} \qquad \Sigma'_D = \{r(g'), w(g') : g' \in G'\}\,.$$

The states of $\mathcal{C}'$ are

$$S' = S \cup \{[s, a, s'] : s, s' \in S, \text{ and } a = \overline{\mathtt{r}}(g) \text{ or } a = \overline{\mathtt{w}}(g) \text{ for some } g \in G\}\,.$$

The new states of the form $[s, a, s']$ represent the situation when the contributor has declared that he wants to do the transition $s \xrightarrow{a} s'$ and move to s'. In order to really move to s', he needs to wait for a confirmation from the leader that the action a has been taken into account. This mechanism is captured by the following transitions of $\mathcal{C}'$:

$$s \xrightarrow{\overline{w}(\overline{\mathtt{w?}}(g))} [s, \overline{\mathtt{w}}(g), s'] \xrightarrow{\overline{r}(\overline{\mathtt{w}}(g))} s' \qquad \text{if} \quad s \xrightarrow{\overline{w}(g)} s' \text{ in } \mathcal{C}$$
$$s \xrightarrow{\overline{w}(\overline{\mathtt{r?}}(g))} [s, \overline{\mathtt{r}}(g), s'] \xrightarrow{\overline{r}(\overline{\mathtt{r}}(g))} s' \qquad \text{if} \quad s \xrightarrow{\overline{r}(g)} s' \text{ in } \mathcal{C}$$

The states of $\mathcal{D}'$ are:

$$T' = \{[t, x] : t \in T, \text{ and } x = g,\ x = \overline{\mathtt{w}}(g), \text{ or } x = \overline{\mathtt{r}}(g), \text{ for some } g \in G\}\,,$$

where the component x is supposed to store the value of the register of the $(\mathcal{C}, \mathcal{D})$-system being simulated. It can also be a read or write operation, when $\mathcal{D}'$ is in the process of confirming a contributor operation. The transitions of $\mathcal{D}'$ are:

$$[t, g] \xrightarrow{w(\mathtt{w}(h))} [t', h] \quad \text{if} \quad t \xrightarrow{w(h)} t' \text{ in } \mathcal{D}$$
$$[t, g] \xrightarrow{w(\mathtt{r}(g))} [t', g] \quad \text{if} \quad t \xrightarrow{r(g)} t' \text{ in } \mathcal{D}$$
$$[t, g] \xrightarrow{r(\overline{\mathtt{w?}}(h))} [t, \overline{\mathtt{w}}(h)] \xrightarrow{w(\overline{\mathtt{w}}(h))} [t, h] \quad \text{for all } t, g, h$$
$$[t, g] \xrightarrow{r(\overline{\mathtt{r?}}(g))} [t, \overline{\mathtt{r}}(g)] \xrightarrow{w(\overline{\mathtt{r}}(g))} [t, g] \quad \text{for all } t, g\,.$$

So transitions of $\mathcal{D}$ are simply reflected by transitions of $\mathcal{D}'$: the value of the register of the $(\mathcal{C}, \mathcal{D})$-system is stored in the state of $\mathcal{D}'$, and the operation being performed is written into the register of $\mathcal{D}'$. When a request of an

operation from a contributor is read then it is performed on the value stored in the state and the confirmation of this operation is written into the register of the $(\mathcal{C}', \mathcal{D}')$-system.

In order to state the correspondence between traces of the $(\mathcal{C}, \mathcal{D})$-system and that of the $(\mathcal{C}', \mathcal{D}')$-system we define some operations on traces. The first one transforms a sequence over the alphabet of the $(\mathcal{C}', \mathcal{D}')$-system into a sequence over the alphabet of the $(\mathcal{C}, \mathcal{D})$-system. A sequence $\mathit{trans}(u)$ is obtained from u by:

1. removing all operations of contributors and all read operations of the leader, and

2. replacing all write operations $w(a)$ of the leader by a, for example $w(\overline{\mathtt{r}}(g))$ is replaced by $\overline{r}(g)$.

So the operation $\mathit{trans}(u)$ is the sequence of operations that are written by $\mathcal{D}'$ into the register.

The second operation uses *stuttering expansions* of sequences over the alphabet of $(\mathcal{C}, \mathcal{D})$-systems wrt. contributor actions. Let $u = u_0 a_0 u_1 a_1 u_2 \cdots$ be a finite or infinite word over $\Sigma_C \cup \Sigma_D$, with $a_i \in \Sigma_C$ and $u_i \in \Sigma_D^*$ for all i (or $u_i \in \Sigma_D^\omega$ if u is infinite but the sequence $u_0, u_1, \ldots$ is finite and of length i). We write $v \in \mathit{stutt}(u)$ if there exists a function $f : \mathbb{N} \to \mathbb{N}^+$ such that $v = u_0 a_0^{f(0)} u_1 a_1^{f(1)} u_2 \cdots$. This operation is required because a single confirmation by the leader of a contributor's request can satisfy several identical requests.

Proposition 34 Let $(\mathcal{C}', \mathcal{D}')$ be obtained from a $(\mathcal{C}, \mathcal{D})$-system as described above. If u is a trace of the $(\mathcal{C}, \mathcal{D})$-system then there is a trace u' of the $(\mathcal{C}', \mathcal{D}')$-system such that $\mathit{trans}(u') = u$. If u' is a trace of $(\mathcal{C}', \mathcal{D}')$-system then there is some $u \in \mathit{stutt}(\mathit{trans}(u'))$ that is a trace of the $(\mathcal{C}, \mathcal{D})$-system.

The proposition above, and then Theorem 5, will follow from the next lemmas.

Lemma 35 If u is a trace of $(\mathcal{C}, \mathcal{D})$ then the trace u' obtained by replacing

$$\begin{array}{rcl}
w(g) & \text{by} & w(\mathtt{w}(g)) \\
r(g) & \text{by} & w(\mathtt{r}(g)) \\
\overline{w}(g) & \text{by} & \overline{w}(\overline{\mathtt{w?}}(g))\, r(\overline{\mathtt{w?}}(g))\, w(\overline{\mathtt{w}}(g))\, \overline{r}(\overline{\mathtt{w}}(g)) \\
\overline{r}(g) & \text{by} & \overline{w}(\overline{\mathtt{r?}}(g))\, r(\overline{\mathtt{r?}}(g))\, w(\overline{\mathtt{r}}(g))\, \overline{r}(\overline{\mathtt{r}}(g))
\end{array}$$

is a trace of $(\mathcal{C}', \mathcal{D}')$. If u is finite and there exists a run of the $(\mathcal{C}, \mathcal{D})$-system over u ending in (M, t, g), then there exists a run of the $(\mathcal{C}', \mathcal{D}')$-system over u' ending in $(M, [t, g], a)$, where a is the last action of u.

Lemma 36 Let $(n[s_{init}], [t_{init}, g_{init}], g'_{init}) \xrightarrow{u'} (M', [t, x], g')$ be a run of the $(\mathcal{C}', \mathcal{D}')$-system. We can construct by induction on the length of u' a run $(n[s_{init}], t_{init}, g_{init}) \xrightarrow{u} (M, t, g)$ of the $(\mathcal{C}, \mathcal{D})$-system and a multiset $N \leq M'$ such that:

1. $u \in \mathit{stutt}(\mathit{trans}(u'))$.
2. If $x \in G$ then $x = g$. If x is of the form $\overline{\mathtt{r}}(h)$ then $h = g$.
3. For all $s \in S$,
$$M(s) = M'(s) + \sum_{a \in \{\overline{r}, \overline{w}\} \times G} \sum_{s' \in S} M'([s, a, s']) - N([s, a, s']) + N([s', a, s])\,.$$
4. If $g' = \overline{\mathtt{w?}}(h)$ or $x = \overline{\mathtt{w}}(h)$, then we have the strict inequality
$$\sum_{s, s' \in S} M'([s, \overline{\mathtt{w}}(h), s']) > \sum_{s, s' \in S} N([s, \overline{\mathtt{w}}(h), s'])\,.$$
Similarly if $g' = \overline{\mathtt{r?}}(h)$ or $x = \overline{\mathtt{r}}(h)$.
5. If $g' = \overline{\mathtt{w}}(h)$, then $N([s, \overline{\mathtt{w}}(h), s']) = M'([s, \overline{\mathtt{w}}(h), s'])$ for every $s, s' \in S$. Similarly if $g' = \overline{\mathtt{r}}(g)$.

Let us explain the role of the multi-set N. States of the form $[s, a, s']$ can be thought of as transitory states of $\mathcal{C}$. They are counted as s, except for those in N that are counted as s'. Conditions 4 and 5 can be interpreted as follows: between actions $r(\overline{\mathtt{w?}}(h))$ and $w(\overline{\mathtt{w}}(h))$, there is at least one contributor in a state $[s, \overline{\mathtt{w}}(h), s']$ that is counted as being in s; after $w(\overline{\mathtt{w}}(h))$, all contributors in $[s, \overline{\mathtt{w}}(h), s']$ are counted as in s'.

Proof. The proof is by induction on the length of u'. We have cases depending on the last action. Given a transition
$$(M'_1, [t_1, x_1], g'_1) \xrightarrow{a'} (M'_2, [t_2, x_2], g'_2) \quad \text{in} \quad (\mathcal{C}', \mathcal{D}')$$
and $N_1 \leq M'_1$, (M_1, t_1, g_1) satisfying the invariants, we define (M_2, t_2, g_2) and $N_2 \leq M'_2$ satifying the invariants and such that
$$(M_1, t_1, g_1) \xrightarrow{a} (M_2, t_2, g_2) \quad \text{in} \quad (\mathcal{C}', \mathcal{D}')$$
where $a \in \mathit{stutt}(\mathit{trans}(a'))$ (we may have $a = \varepsilon$). We first consider all the possible actions of $\mathcal{D}'$, and later those of $\mathcal{C}'$. When not stated otherwise, we keep $N_2 = N_1$.

- $(M', [t_1, g], *) \xrightarrow{w(\mathtt{w}(h))} (M', [t_2, h], \mathtt{w}(h))$ is simulated by
$(M, t_1, g) \xrightarrow{w(h)} (M, t_2, h)$

- $(M', [t_1, g], *) \xrightarrow{w(\mathtt{r}(g))} (M', [t_2, g], \mathtt{r}(g))$ is simulated by
  $(M, t_1, g) \xrightarrow{r(g)} (M, t_2, g)$

- $(M', [t, g], \overline{\mathtt{w?}}(h)) \xrightarrow{r(\overline{\mathtt{w?}}(h))} (M', [t, \overline{\mathtt{w}}(h)], \overline{\mathtt{w?}}(h))$ is simulated by no action.

- $(M', [t, g], \overline{\mathtt{r?}}(g)) \xrightarrow{r(\overline{\mathtt{r?}}(g))} (M', [t, \overline{\mathtt{r}}(g)], \overline{\mathtt{r?}}(g))$ is simulated by no action.

- $(M', [t, \overline{\mathtt{w}}(h)], *) \xrightarrow{w(\overline{\mathtt{w}}(h))} (M', [t, h], \overline{\mathtt{w}}(h))$ is simulated by
  $(M_1, t, g) \xrightarrow{\overline{w}(h)^k} (M_2, t, h)$

  Intuitively in this step we do all the writes that waited to be done, so:

  (i) $k = \sum_{s,s' \in S} d_{s,s'}$ with $d_{s,s'} = M'[s, \overline{\mathtt{w}}(h), s'] - N_1[s, \overline{\mathtt{w}}(h), s']$; this number says how many contributors there are in state s that want to do $\overline{\mathtt{w}}(h)$ and go to s'. Note that by invariant 4, $k > 0$.

  (ii) $M_2(s) = M_1(s) - \sum_{s' \in S} d_{s,s'} + \sum_{s' \in S} d_{s',s}$.

  (iii) N is updated from N_1 to N_2 with $N_2[s, \overline{\mathtt{w}}(h), s'] = M'[s, \overline{\mathtt{w}}(h), s']$ for every $s, s' \in S$, and $N_2[s, a, s'] = N_1[s, a, s']$ if $a \neq \overline{\mathtt{w}}(h)$.

  The run is well-defined since by invariant 3, $M_1(s) \geq \sum_{s' \in S} d_{s,s'}$. Invariant 5 is preserved thanks to item *(iii)*. A small calculation shows that invariant 3 is preserved:

$$\begin{aligned}
M_2(s) &= M'(s) + \sum_{a \in \{\overline{r}, \overline{w}\} \times G} \sum_{s' \in S} M'[s, a, s'] - N_1[s, a, s'] + N_1[s', a, s] \\
&\quad - \sum_{s' \in S'} d_{s,s'} + \sum_{s' \in S'} d_{s',s} \\
&= M'(s) + \sum_{a \neq \overline{\mathtt{w}}(h)} \sum_{s' \in S} M'[s, a, s'] - N_1[s, a, s'] + N_1[s', a, s] \\
&\quad + \sum_{s' \in S} M'([s', \overline{\mathtt{w}}(h), s]) \\
&= M'(s) + \sum_{a \in \{\overline{r}, \overline{w}\} \times G} \sum_{s' \in S} M'[s, a, s'] - N_2[s, a, s'] + N_2[s', a, s] \,.
\end{aligned}$$

- $(M', [t, \overline{\mathtt{r}}(g)], *) \xrightarrow{w(\overline{\mathtt{r}}(g))} (M', [t, h], \overline{\mathtt{r}}(g))$ is simulated by
  $(M_1, t, g) \xrightarrow{\overline{r}(g)^k} (M_2, t, g)$

  where k and the multi-sets are defined as for writes.

- $(M_1', [t,x], *) \xrightarrow{\overline{w}(\overline{\mathtt{w?}}(h))} (M_2', [t,x], \overline{\mathtt{w?}}(h))$ is simulated by no operation from (M,t,g).

  We have $M_2' = M_1' - [s] + [[s, \overline{\mathtt{w}}(g), s']]$ for some s, s'. Invariant 3 is preserved, as M and N are not modified and:

  $$M_1'(s) + M_1'([s, \overline{\mathtt{w}}(g), s']) = M_2'(s) + M_2'([s, \overline{\mathtt{w}}(g), s'])$$

  Moreover, we have $N \leq M_2'$ and

  $$M_2'([s, \overline{\mathtt{w}}(g), s']) > M_1'([s, \overline{\mathtt{w}}(g), s']) \geq N([s, \overline{\mathtt{w}}(g), s'])$$

  so invariant 4 is verified.

- $(M_1', [t,x], *) \xrightarrow{\overline{w}(\overline{\mathtt{r?}}(h))} (M_2', [t,x], \overline{\mathtt{r?}}(h))$ is simulated by no operation from (M,t,g). The invariant is preserved, as in the previous case.

- $(M_1', [t,x], \overline{\mathtt{w}}(h)) \xrightarrow{\overline{r}(\overline{\mathtt{w}}(h))} (M_2', [t,x], \overline{\mathtt{w}}(h))$ is simulated by no operation; however, we need to update N_1 to keep the invariant. For this we take the state $[s, \overline{\mathtt{w}}(h), s']$ that changed to s' while going from M_1' to M_2', and obtain N_2 by substracting 1 from $N_1([s, \overline{\mathtt{w}}(h), s'])$. This is possible since by invariant 5, $N_1([s, \overline{\mathtt{w}}(h), s']) = M_1'([s, \overline{\mathtt{w}}(h), s'])$.

  Clearly, invariant 5 is preserved. For invariant 3, observe that:

  $$M_1'([s, \overline{\mathtt{w}}(h), s']) - N_1([s, \overline{\mathtt{w}}(h), s']) = M_2'([s, \overline{\mathtt{w}}(h), s']) - N_2([s, \overline{\mathtt{w}}(h), s'])$$
  $$M_1'(s') + N_2([s, \overline{\mathtt{w}}(h) s']) = M_1'(s') + N_2([s, \overline{\mathtt{w}}(h) s'])$$

  so the equalities at s and s' are still true. The others do not change.

- $(M_1', [t,g], \overline{\mathtt{r}}(g)) \xrightarrow{\overline{r}(\overline{\mathtt{r}}(g))} (M_2', [t,g], \overline{\mathtt{r}}(g))$ is simulated similarly.

□

Unfortunately, the above construction does not preserve run maximality. First, it introduces deadlocks: a contributor can declare that he wants to do a read, but this read turns out to be illegal, or is simply ignored by the leader; as a result the contributor gets stuck in state $[s, \overline{\mathtt{r}}(g), s']$. Moreover, a deadlock $([s_1, \ldots, s_n], t, g)$ in $(\mathcal{C}, \mathcal{D})$ does not correspond to a deadlock $([s_1, \ldots, s_n], [t,g], g')$ in $(\mathcal{C}', \mathcal{D}')$, since a contributor may be unable to execute an action $s \xrightarrow{\overline{r}(h)} s'$ in $(\mathcal{C}, \mathcal{D})$, but is always able to execute $s \xrightarrow{\overline{w}(\overline{\mathtt{r?}}(h))} [s, \overline{\mathtt{r}}(h), s']$ in $(\mathcal{C}', \mathcal{D}')$.

We can modify the $(\mathcal{C}', \mathcal{D}')$-system in order to have a correspondence between maximal runs of the new system and maximal runs of the $(\mathcal{C}, \mathcal{D})$-system.

Lemma 37 There is a $(\mathcal{C}'', \mathcal{D}'')$-system and a register value # such that:

1. If u is a (maximal) trace of the $(\mathcal{C}, \mathcal{D})$-system, then there exists a (maximal) trace v of the $(\mathcal{C}'', \mathcal{D}'')$-system with no occurence of $w(\#)$ and such that $\mathit{trans}(v) = u$.

2. If v is a (maximal) trace of the $(\mathcal{C}'', \mathcal{D}'')$-system with no occurence of $w(\#)$, then there exists a (maximal) trace u of the $(\mathcal{C}, \mathcal{D})$-system such that $u \in \mathit{stutt}(\mathit{trans}(v))$.

Proof. The result already holds for infinite runs with $(\mathcal{C}', \mathcal{D}')$. The idea is to modify the $(\mathcal{C}', \mathcal{D}')$-system so that the leader can guess the end of a finite run corresponding to a maximal run of the $(\mathcal{C}, \mathcal{D})$-system. Whenever her guess is wrong, she can detect it, and write some special value $\#$ to the register.

We start by definying $\mathcal{D}''$. For a value $g \in G$, we call a state t of $\mathcal{D}$ a *g-deadlock* if from t there is no outgoing transition labeled by $r(g)$ or by a write of some value. The leader $\mathcal{D}''$ is obtained by adding states t_f, $t_\#$ and $t'_\#$ to $\mathcal{D}'$, for all states t of $\mathcal{D}$, and transitions

$$\begin{aligned} [t, g] &\xrightarrow{w(g)} t_f && \text{for all } g \in G, \text{ and } g\text{-deadlock states } t \text{ of } \mathcal{D} \\ t_f &\xrightarrow{r(a)} t_\# \xrightarrow{w(\#)} t'_\# && \text{for all } a \in \{\#\} \cup \{\overline{\texttt{r?}}(g), \overline{\texttt{w?}}(g) \mid g \in G\} \end{aligned}$$

Similarly, $\mathcal{C}''$ is obtained by adding to $\mathcal{C}'$ states s_f, $s_\#$, $s'_\#$, for every $s \in S$, and transitions:

$$\begin{aligned} [s, a, s'] &\xrightarrow{\overline{r}(a)} s'_f && \text{if } \mathcal{C} \text{ cannot do a write from } s \\ s_f &\xrightarrow{\overline{r}(g)} s_\# && \text{if } \mathcal{C} \text{ can do } \overline{r}(g) \text{ from } s \\ [s, a, s'] &\xrightarrow{\overline{r}(g)} s_\# && \text{for all } g \in G \\ s &\xrightarrow{\overline{r}(g)} s_\# && \text{for all } g \in G \\ s_\# &\xrightarrow{\overline{w}(\#)} s'_\# \end{aligned}$$

Intuitively, an error $\#$ can occur in two situations. The first is when the leader has guessed the end of the run by moving to a state t_f, but some contributor hasn't: he is not in a state of the form s_f. The second is when the guessed end configuration $([s^1_f, \ldots, s^n_f], t_f, g)$ does not correspond to a deadlock, i.e., $\mathcal{C}$ can read g from some state s^i.

Observe that the $(\mathcal{C}'', \mathcal{D}'')$-system has the same set of infinite traces as the $(\mathcal{C}', \mathcal{D}')$-system: in an infinite run, states $t_f, t_\#, t'_\#$, and thus also $s_\#, s'_\#$ are never reached, and states s_f can be replaced by s.

From a finite maximal run of the $(\mathcal{C}, \mathcal{D})$-system, we can construct a finite maximal run of the $(\mathcal{C}'', \mathcal{D}'')$-system with no occurence of $w(\#)$, as in Lemma 35, but having the contributors move to states s_f instead of s in their last transitions, and adding a transition $[t, g] \xrightarrow{w(g)} t_f$ at the end of the run.

Conversely, consider a finite maximal run with labeling v of the $(\mathcal{C}'', \mathcal{D}'')$-system, with no occurence of $w(\#)$. The last action of the leader must be some $[t, g] \xrightarrow{w(g)} t_f$, because she can always do a write from a state of the form $[t, g]$. Due to the transitions we have added in $\mathcal{D}'$, the value of the register in the final configuration cannot be $\overline{\texttt{r?}}(h)$, $\overline{\texttt{w?}}(h)$ nor $\#$. The value is then g because the last action of the leader was $w(g)$ and contributors can write only $\overline{\texttt{r?}}(h)$, $\overline{\texttt{w?}}(h)$ or $\#$. Due to transitions added above, none of the contributors in the $(\mathcal{C}'', \mathcal{D}'')$-system can be in a state of the form $s \in S$ or $[s, a, s']$. So all the contributors are in states of the form s_f such that s is a g-deadlock. By construction, writing $v = v_1 w(g) v_2$ and $([t_f^1, \ldots, t_f^n], t_f, g)$ the end configuration, there is also a run $(n[s_{init}], [t_{init}, g_{init}], g_{init}) \xrightarrow{v_1 v_2} ([s^1, \ldots, s^n], [t, g], g')$ in the $(\mathcal{C}', \mathcal{D}')$-system. By Lemma 36, this leads a run $(n[s_{init}], t_{init}, g_{init}) \xrightarrow{u} ([s^1, \ldots, s^n], t, g)$ in the $(\mathcal{C}, \mathcal{D})$-system, for some $u \in \mathit{stutt}(\mathit{trans}(v))$. It is maximal since from none of t, $s^1, \ldots, s^n$ it is possible to do a write or a read of g. □

We can now prove Theorem 5.

Proof of Thm. 5. For $\tilde{\mathcal{P}}$ we take the property

$$\tilde{\mathcal{P}} = \{u'' : \mathit{trans}(u'') \in \mathcal{P} \text{ and there is no } w(\#) \text{ in } u''\}$$

If $\mathcal{P}$ is regular, so is $\tilde{\mathcal{P}}$. Similarly if $\mathcal{P}$ is defined by an LTL formula.

Let $(\tilde{\mathcal{C}}, \tilde{\mathcal{D}}) := (\mathcal{C}'', \mathcal{D}'')$ as in Lemma 37. If the $(\mathcal{C}, \mathcal{D})$-system has a maximal trace $u \in \mathcal{P}$, by Lemma 37, the $(\tilde{\mathcal{C}}, \tilde{\mathcal{D}})$-system has a maximal trace u'' with no occurence of $w(\#)$ and such that $\mathit{trans}(u'') = u$.

If the $(\tilde{\mathcal{C}}, \tilde{\mathcal{D}})$-system has a maximal trace u'' such that $u'' \in \tilde{\mathcal{P}}$, then by Lemma 37 and since u'' contains no occurrence of $w(\#)$, there exists a maximal trace u in the $(\mathcal{C}, \mathcal{D})$-system such that $u \in \mathit{stutt}(\mathit{trans}(u''))$. We have $\mathit{trans}(u'') \in \mathcal{P}$, by the definition of $\tilde{\mathcal{P}}$. Since $\mathcal{P}$ is $\mathcal{C}$-expanding, u is also in $\mathcal{P}$. □

A further byproduct of the simulation technique used in Theorem 5 is that we can simulate a $(\mathcal{C}, \mathcal{D})$-system with m shared registers by one with a single shared register:

Theorem 38 *Let m be fixed. For every $(\mathcal{C}, \mathcal{D})$-system with m registers, and every $\mathcal{C}$-expanding regular (resp. LTL) property $\mathcal{P} \subseteq (\Sigma_C \cup \Sigma_D)^\infty$, there exists a $(\tilde{\mathcal{C}}, \tilde{\mathcal{D}})$-system with one register and a regular (resp. LTL) property $\tilde{\mathcal{P}} \subseteq (\widetilde{\Sigma}_D)^\infty$, where $\widetilde{\Sigma}_D$ is the input alphabet of $\tilde{\mathcal{D}}$, such that:*

> *the $(\mathcal{C}, \mathcal{D})$-system has a maximal trace $u \in \mathcal{P}$ iff the $(\tilde{\mathcal{C}}, \tilde{\mathcal{D}})$-system has a maximal trace u' whose projection on $\widetilde{\Sigma}_D$ is in $\tilde{\mathcal{P}}$.*

Moreover, the $(\tilde{\mathcal{C}}, \tilde{\mathcal{D}})$-system and the property $\tilde{\mathcal{P}}$ are effectively computable in polynomial time.

8 Conclusion

We studied verification questions for parametrized, asynchronous, shared-memory pushdown systems consisting of a leader process and arbitrarily many, anonymous contributor processes, as in Hague's model [15, 12]. First, we answered an open question of [10], by showing that the complexity of checking liveness in this model is PSPACE-complete. Then we established the complexity of checking universal reachability. The developed techniques allowed us to consider a more general problem, of verifying regular properties that refer to both leader and contributors, but are stutter-invariant w.r.t. contributor actions. We have shown that this problem is decidable and NEXPTIME-complete.

References

[1] M. F. Atig, A. Bouajjani, and S. Qadeer. Context-bounded analysis for concurrent programs with dynamic creation of threads. *Logical Methods in Computer Science*, 7(4):1–48, 2011.

[2] T. Ball, S. Chaki, and S. K. Rajamani. Parameterized verification of multithreaded software libraries. In *TACAS'01*, LNCS, pages 158–173. Springer, 2001.

[3] R. Bloem, S. Jacobs, A. Khalimov, I. Konnov, S. Rubin, H. Veith, and J. Widder. *Decidability of Parameterized Verification.* Morgan & Claypool Publishers, 2015.

[4] A. Bouajjani, J. Esparza, and O. Maler. Reachability analysis of pushdown automata: Application to model-checking. In *CONCUR '97*, volume 1243 of *LNCS*, pages 135–150. Springer, 1997.

[5] A. Bouajjani, J. Esparza, S. Schwoon, and J. Strejcek. Reachability analysis of multithreaded software with asynchronous communication. In *FSTTCS'05*, LNCS, pages 348–359. Springer, 2005.

[6] P. Bouyer, N. Markey, M. Randour, A. Sangnier, and D. Stan. Reachability in networks of register protocols under stochastic schedulers. In *ICALP'16*, LIPIcs. Leibniz-Zentrum für Informatik, 2016. To appear.

[7] E. M. Clarke, O. Grumberg, and S. Jha. Verifying parameterized networks. *ACM Trans. Program. Lang. Syst.*, 19(5):726–750, 1997.

[8] B. Courcelle. On constructing obstruction sets of words. *Bulletin of EATCS*, 1991.

[9] G. Delzanno. Parameterized verification and model checking for distributed broadcast protocols. In *ICGT'14*, LNCS, pages 1–16. Springer, 2014.

[10] A. Durand-Gasselin, J. Esparza, P. Ganty, and R. Majumdar. Model checking parameterized asynchronous shared-memory systems. In *CAV*, 2015.

[11] J. Esparza, A. Finkel, and R. Mayr. On the verification of broadcast protocols. In *LICS'99*, pages 352–359. IEEE, 1999.

[12] J. Esparza, P. Ganty, and R. Majumdar. Parameterized verification of asynchronous shared-memory systems. *J. ACM*, 63(1):10, 2016.

[13] K. Etessami. A note on a question of Peled and Wilke regarding stutter-invariant LTL. *Inf. Process. Lett.*, 75(6):261–263, 2000.

[14] S. A. German and P. A. Sistla. Reasoning about systems with many processes. *J. ACM*, 39(3):675–735, 1992.

[15] M. Hague. Parameterised pushdown systems with non-atomic writes. In *FSTTCS'11*, LIPIcs, pages 457–468. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2011.

[16] V. Kahlon. Parameterization as abstraction: A tractable approach to the dataflow analysis of concurrent programs. In *LICS*, 2008.

[17] A. Kaiser, D. Kroening, and T. Wahl. Dynamic cutoff detection in parameterized concurrent programs. In *CAV'10*, LNCS, pages 645–659. Springer, 2010.

[18] Y. Kesten, A. Pnueli, E. Shahar, and L. D. Zuck. Network invariants in action. In *CONCUR'02*, LNCS, pages 101–115. Springer, 2002.

[19] S. La Torre, P. Madhusudan, and G. Parlato. Model-checking parameterized concurrent programs using linear interfaces. In *CAV'10*, LNCS, pages 629–644. Springer, 2010.

[20] S. La Torre, P. Madhusudan, and G. Parlato. Sequentializing parameterized programs. In *FIT'12*, volume 87 of *EPTCS*, pages 34–47, 2012.

[21] S. La Torre, A. Muscholl, and I. Walukiewicz. Safety of parametrized asynchronous shared-memory systems is almost always decidable. In *CONCUR'15*, volume 42 of *LIPIcs*, pages 72–84. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.

[22] K. S. Namjoshi and R. J. Trefler. Analysis of dynamic process networks. In *TACAS'15*, LNCS, pages 164–178. Springer, 2015.

[23] D. A. Peled and T. Wilke. Stutter-invariant temporal properties are expressible without the next-time operator. *Inf. Process. Lett.*, 63(5):243–246, 1997.

[24] G. Ramalingam. Context-sensitive synchronization-sensitive analysis is undecidable. *ACM Trans. Program. Lang. Syst. (TOPLAS)*, 22(2):416–430, 2000.

A Proof of Lemma 7

Sketch of proof. The idea is to "distribute" a run of a contributor into several runs with smaller stacks.

Consider a finite run ρ of some copy of $\mathcal{C}$ in which the effective stack-height goes above N. Let s be one of the configurations in the run with effective stack-height greater than N. We write $A_1 \cdots A_n A_{n+1} A_{n+2}$ its effective stack, and $A_1 \cdots A_n \alpha$ its total stack. All symbols of the effective stack, except possibly A_{n+2}, are eventually popped in the run. In particular, $A_1, \ldots, A_n$ are popped strictly before the last action in the run. We consider the positions just after the symbols $A_1, \ldots, A_n$ are last pushed before reaching configuration s (resp. popped after s). So we have:

$$\begin{aligned}\rho : (p_{init} A^C_{init}) &\xrightarrow{u_n} (p_n A_n \alpha) \xrightarrow{u_{n-1}} (p_{n-1} A_{n-1} A_n \alpha) \xrightarrow{u_{n-2}} \\ &\cdots \xrightarrow{u_1} (p_1 A_1 \cdots A_n \alpha) \xrightarrow{v_1} (r_1 A_2 \cdots A_n \alpha) \xrightarrow{v_2} \cdots \xrightarrow{v_n} (r_n \alpha) \xrightarrow{v} (p_f \alpha_f)\end{aligned}$$

where $u_1, \ldots, u_n, v_1, \ldots, v_n, v \in \Sigma_C^+$, and in the part $(p_i A_i \cdots A_n \alpha) \xrightarrow{u_{i+1} \cdots u_1} (p_1 A_1 \cdots A_n \alpha) \xrightarrow{v_1 \cdots v_i} (r_i A_{i+1} \cdots A_n)$, the bottom $A_i \cdots A_n \alpha$ of the stack is never modified except by the last action, which pops A_i.

Since $n > 2|P|^2|\Gamma_C|$, there must be three indices $1 \leq i < j < k \leq n$ such that

$$(p_i, A_i, r_i) = (p_j, A_j, r_j) = (p_k, A_k, r_k)\,.$$

We can then construct two smaller runs ρ_1 and ρ_2 of $\mathcal{C}$, by removing respectively the parts $u_{j-1} \cdots u_i$ and $v_{i+1} \cdots v_j$, or $u_{k-1} \cdots u_j$ and $v_{j+1} \cdots v_k$, from ρ. That is,

$$\begin{aligned}\rho_1 : (p_{init} A^C_{init}) &\xrightarrow{u_n \cdots u_j} (p_j A_j \cdots A_n \alpha) \xrightarrow{u_{i-1} \cdots u_1} (p_1 A_1 \cdots A_i A_{j+1} \cdots A_n \alpha) \\ &\xrightarrow{v_1 \cdots v_i} (r_j A_{j+1} \cdots A_n \alpha) \xrightarrow{v_{j+1} \cdots v_n v} (p_f \alpha_f)\end{aligned}$$

and

$$\begin{aligned}\rho_2 : (p_{init} A^C_{init}) &\xrightarrow{u_n \cdots u_k} (p_k A_k \cdots A_n \alpha) \xrightarrow{u_{j-1} \cdots u_1} (p_1 A_1 \cdots A_j A_{k+1} \cdots A_n \alpha) \\ &\xrightarrow{v_1 \cdots v_j} (r_j A_{k+1} \cdots A_n \alpha) \xrightarrow{v_{k+1} \cdots v_n v} (p_f \alpha_f)\,.\end{aligned}$$

Observe that

- any transition in ρ can always be associated with a transition of ρ_1, a transition of ρ_2, or both.
- since $v \neq \varepsilon$, both ρ_1 and ρ_2 end in the same configuration as ρ.

In a run of a $(\mathcal{C}, \mathcal{D})$-system, a contributor executing ρ can thus be replaced by two contributors executing respectively ρ_1 and ρ_2. They progress together

as in the original run on the common parts, and for the other parts, one of the two contributor waits while the other performs the actions of the original run. We repeat this until no contributor ever uses an effective stack of height greater than N: at each step, we replace one contributor by two contributors, but performing strictly shorter runs, so this procedure terminates.

B Proof of Lemma 11

For the left to right direction, we apply Lemma 9 to obtain a run of the $(\mathcal{C}_{fin}, \mathcal{D})$-system of the form $(M, qA, g) \xrightarrow{u} (M, qA\alpha', g)$, where $(M, qA\alpha, g)$ is reachable for some α, and u is of the form $u = u'r(g)$ or $u = u'w(g)$, containing some occurrence of $\top$.

We write $M = [p_1, \ldots, p_n]$. There are words $v_0 \in \Sigma_D^*$, $v_1, \ldots, v_n \in \Sigma_C^*$ such that

$$u \in v_0 \sqcup\!\sqcup v_1 \sqcup\!\sqcup \cdots \sqcup\!\sqcup v_n \quad (v \text{ is a shuffle of } v_0, \ldots, v_n),$$
$$qA \xrightarrow{v_0} qA\alpha' \text{ in } \mathcal{D}\,,$$

and for some permutation σ of $\{1, \ldots, n\}$, for all $1 \leq i \leq n$,

$$p_i \xrightarrow{v_i} p_{\sigma(i)} \text{ in } \mathcal{C}_{fin}\,.$$

By repeating the run if necessary, we can assume that σ is the identity. More precisely, for k such that σ^k is the identity, we can replace u by $(u)^k$, α' by $(\alpha')^k$, and (v_i) by $(v_i)^k$ for all i.

We let $\overline{u}$ be the word obtained by replacing each first occurrence of $\overline{w}(h)$ in u by $\nu(h)\overline{w}(h)$, and $v = \overline{u}|_{\Sigma_{D,\nu}}$. Note that $last(v) = g$, and $v|_{\Sigma_D} = v_0$. Using the fact that $qA \xrightarrow{v_0} qA\alpha'$ in $\mathcal{D}$, we can show by induction on the length of v that there is a run of the form $(\emptyset, qA, g) \xrightarrow{v} (K, qA\alpha', g)$ in $\mathcal{D}^\kappa$. All we need to show is that each read in v is enabled, i.e. that for all prefix $v'r(h)$ of v, either $last(v') = h$ or $\nu(h)$ occurs in v'. Consider a prefix $v'r(h)$ for which $last(v') \neq h$. Then in the corresponding prefix $\overline{u}'r(h)$ of $\overline{u}$ (i.e. with $\overline{u}'|_{\Sigma_{D,\nu}} = v'$), there must be an occurrence of $\overline{w}(h)$ in $\overline{u}'$, and the first such occurrence is preceded by $\nu(h)$. So $\nu(h)$ also occurs in v'.

Denote by $h_1, \ldots, h_m$ the values h such that $\overline{w}(h)$ occurs in u, ordered according to their first occurrences. For all $1 \leq j \leq m$, we let i_j be the index such that the first occurrence of $\overline{w}(h_j)$ in u comes from v_{i_j}. We now show that v is ω-supported from $(p_{i_1}, \ldots, p_{i_m})$.

For all $1 \leq j \leq m$, we must find a word u^j such that

$$u^j \in \left(\Sigma_{C,D,\nu}^* \Sigma_{D,\nu}\right) \cap \left(\Sigma_{C,D,\nu}^* \nu(h_i)\overline{w}(h_i)\Sigma_{C,D,\nu}^*\right)$$
$$u^j|_{\Sigma_{D,\nu}} = v$$
$$(\emptyset, p_{i_j}, g) \xrightarrow{u^j} (\{h_1, \ldots, h_m\}, p_{i_j}, g) \text{ in } \mathcal{C}_{fin}^\kappa\,.$$

We let u^j be the restriction of $\overline{u}$ to positions coming from v_0, v_{i_j} and positions of the $\nu(h_k)$. By definition, $u^j \in \Sigma^*_{C,D,\nu}\nu(h_i)\overline{w}(h_i)\Sigma^*_{C,D,\nu}$ and $u^j|_{\Sigma_{D,\nu}} = v$. Moreover, since u ends with $w(g)$ or $r(g)$, so does u^j, and $last(u^j) = g$. We show that $(\emptyset, p_{i_j}, g) \xrightarrow{u^j} (\{h_1, \ldots, h_m\}, p_{i_j}, g)$ in $\mathcal{C}^\kappa_{fin}$ similarly to what we did for v, using the fact that $u^j|_{\Sigma_\mathcal{C}} = v_{i_j}$ and $p_{i_j} \xrightarrow{v_{i_j}} p_{i_j}$ in $\mathcal{C}_{fin}$.

For the right to left direction, consider a reachable configuration $(M, qA\alpha, g)$ and a word $v = v_1\nu(h_1)\cdots v_m\nu(h_m)v_{m+1}$ satisfying conditions (1) and (2) of Lemma 11, and let ρ be the run in condition (1). For all i, there exists a word

$$u^i = u^i_1\nu(h_1)\cdots u^i_i\nu(h_i)\overline{w}(h_i)\cdots u^i_m\nu(h_m)u^i_{m+1}$$

such that $u^i_k|_{\Sigma_{\mathcal{D},\nu}} = v_k$, and $\mathcal{C}^\kappa_{fin}$ has a run ρ^i of the form

$$(\emptyset, p_i, g) \xrightarrow{u^i_1\nu(h_1)\cdots u^i_i\nu(h_i)} (\{h_1, \ldots, h_i\}, p'_i, h_i) \xrightarrow{\overline{w}(h_i)} (\{h_1, \ldots, h_i\}, p''_i, h_i) \xrightarrow{u^i_{i+1}\nu(h_{i+1})\ldots u^i_{m+1}} (\{h_1, \ldots, h_m\}, p_i, g)\,.$$

We denote by n_i the total number of capacity reads of value h_i occurring either in ρ or in one of the ρ^j. We show the following lemmas, that describe how to construct an ultimately periodic run of the $(\mathcal{C}_{fin}, \mathcal{D})$-system. The proofs of these lemmas ressemble the one of Lemma 4 in [21], but are more involved since we are not only interested in reachability and we need to track precisely the number of contributors in a given state.

Lemma 39 There is a run of the $(\mathcal{C}_{fin}, \mathcal{D})$-system of the form

$$(\textstyle\sum_{i=1}^m (n_i+1)[p_i], qA, g) \xrightarrow{*} (\sum_{i=1}^m [p_i] + n_i[p''_i], qA\alpha', g)\ .$$

Lemma 40 There is a run of the $(\mathcal{C}_{fin}, \mathcal{D})$-system of the form

$$(\textstyle\sum_{i=1}^m ((n_i+1)[p_i] + n_i[p''_i])\,, qA, g) \xrightarrow{u} (\sum_{i=1}^m (n_i+1)[p_i] + n_i[p''_i], qA\alpha', g)$$

for some $u \in \Sigma^*_{\mathcal{C},\mathcal{D}}\top\Sigma^*_{\mathcal{C},\mathcal{D}}$.

Lemmas 39 and 40 show that there is a Büchi run starting from $(M', qA\alpha, g)$ in the $(\mathcal{C}_{fin}, \mathcal{D})$-system, for any $M' \geq \sum_{i=1}^m (2n_i+1)[p_i]$. Since $(M, qA\alpha, g)$ is reachable in the $(\mathcal{C}_{fin}, \mathcal{D})$-system and $M \geq [p_1, \ldots, p_n]$, by dupplicating the runs of the contributors ending in $p_1, \ldots, p_n$, we obtain that the configuration $(M', qA\alpha, g)$ where $M' = M + \sum_{i=1}^m 2n_i[p_i]$ is also reachable. Hence, the $(\mathcal{C}_{fin}, \mathcal{D})$-system has a Büchi run.

Proof of Lemma 39. The run is obtained as follows. The leader behaves as in ρ, and for all i, one of the contributors, synchronized with the leader,

behaves as in ρ^i. We call this contributor the *main copy* of ρ^i. The remaining contributors starting in p_i follow the main copy of ρ^i up to reaching p'_i, and then stop. Then, each time some process needs to read the value h_i, one of the contributors waiting in state p'_i takes the transition $p'_i \xrightarrow{\overline{w}(h_i)} p''_i$. This is defined more precisely below.

First, we introduce some notations. For any $a \in \Sigma_{\mathcal{D},\mathcal{C},\nu}$, and $n \in \mathbb{N}$, we denote by $a^{(n)}$ the word consisting of n a's, and for any word $w = a_1 \cdots a_j$, we let $w^{(n)} = a_1^{(n)} \cdots a_j^{(n)}$. Recall that $v = v_1\nu(h_1)\cdots v_m\nu(h_m)v_{m+1}$ satisfies conditions (1) and (2) of Lemma 11.

For $1 \le k \le m+1$, we write

$$\begin{aligned} v_k &= \boldsymbol{a_{k,1} \cdots a_{k,\ell_k}}\,, \text{ and} \\ u_k^i &= x_{k,1}^i \boldsymbol{a_{k,1}} \cdots x_{k,\ell_k}^i \boldsymbol{a_{k,\ell_k}} x_{k,\ell_k+1}^i \quad (x_{k,j}^i \in \Sigma_{\mathcal{C}}^*)\,. \end{aligned}$$

We write $x_{k,j}^i$ as $x_{k,j}^i = y_{k,j}^i z_{k,j}^i$, where $y_{k,j}^i$ is the largest prefix of $x_{k,j}^i$ that consists of register reads only. Since a register read is necessarily a read from the initial value or a read from a value written by the leader, it can only follow an action of the leader or another register read. Hence each $z_{k,j}^i$ contains only writes and capacity reads.

We are going to define a trace of the $(\mathcal{C}_{fin}, \mathcal{D})$-system as a shuffle of $v|_{\Sigma_D} = v_1 \cdots v_{m+1}$ and each of the $\big((u_1^i \cdots u_i^i)^{(n_i+1)}\overline{w}(h_i)(u_{i+1}^i \cdots u_{m+1}^i)\big)\,|_{\Sigma_C}$ and $\overline{w}(h_i)^{(n_i)}$. This corresponds to the intuition that until the first $\overline{w}(h_i)$, all (n_i+1) contributors starting in p_i follow the main copy of ρ^i, and then the main copy continues alone. To make sure that all reads are enabled, the trace will be constructed as follows: after each action $\boldsymbol{a_{k,j}}$ of the leader, we put first all register reads that follow it in one of the u^i (i.e., actions from some $y_{k,j+1}^i$), then all writes or capacity reads (i.e., actions from some $z_{k,j+1}^i$).

For all $1 \le k \le m$, define

$$\begin{aligned} w_k = \; & y_{k,1}^1 \cdots y_{k,1}^{k-1}\ (y_{k,1}^k)^{(n_k+1)} \cdots (y_{k,1}^m)^{(n_m+1)} \\ & \qquad z_{k,1}^1 \cdots z_{k,1}^{k-1} (z_{k,1}^k)^{(n_k+1)} \cdots (z_{k,1}^m)^{(n_m+1)} \boldsymbol{a_{k,1}} \\ & \cdots \\ & y_{k,\ell_k}^1 \cdots y_{k,\ell_k}^{k-1} (y_{k,\ell_k}^k)^{(n_k+1)} \cdots (y_{k,\ell_k}^m)^{(n_m+1)} \\ & \qquad z_{k,\ell_k}^1 \cdots z_{k,\ell_k}^{k-1} (z_{k,\ell_k}^k)^{(n_k+1)} \cdots (z_{k,\ell_k}^m)^{(n_m+1)} \boldsymbol{a_{k,\ell_k}} \\ & y_{k,\ell_k+1}^1 \cdots y_{k,\ell_k+1}^{k-1} (y_{k,\ell_k+1}^k)^{(n_k+1)} \cdots (y_{k,\ell_k+1}^m)^{(n_m+1)} \\ & \qquad z_{k,\ell_k}^1 \cdots z_{k,\ell_k}^{k-1} (z_{k,\ell_k+1}^k)^{(n_k+1)} \cdots (z_{k,\ell_k+1}^m)^{(n_m+1)} \overline{w}(h_k)\,, \end{aligned}$$

and similarly w_{m+1}, except we remove the last $\overline{w}(h_k)$. We let $w = w_1 \cdots w_{m+1}$.

We write now $w = b_1 \cdots b_r$, where each b_θ is either one of the $\boldsymbol{a_{k,j}}$, a first occurrence of $\overline{w}(h_i)$, a single letter "a" of one of the $y_{k,j}^i$, $z_{k,j}^i$, or a repetition

$a^{(n_i+1)}$ of a letter "a" in one of the $(y^i_{k,j})^{(n_i+1)}$, $(z^i_{k,j})^{(n_i+1)}$. We define $\overline{b}_\theta$ as follows:

- If $b_\theta = \overline{r}(h)$ and $\overline{r}(h)$ corresponds to a capacity read in some u^i, then $\overline{b}_\theta = \overline{w}(h)\overline{r}(h)$.
- If $b_\theta = \overline{r}(h)^{(n)}$ and $\overline{r}(h)$ corresponds to a capacity read in some u^i, then $\overline{b}_\theta = \overline{w}(h)\overline{r}(h)^{(n)}$.
- If $b_\theta = r(h)$ and $r(h)$ is a capacity read in v, then $\overline{b}_\theta = \overline{w}(h)r(h)$.
- Else, $\overline{b}_\theta = b_\theta$.

We let $\overline{w} = \overline{b}_1 \cdots \overline{b}_r$.

We denote by θ_i the position of the first occurrence of $\overline{w}(h_i)$ in w.

For all $\theta \geq 0$, we let $u^i(\theta)$ be the prefix of u^i associated with $b_1 \cdots b_\theta$, and $(K_\theta, p^i_\theta, g^i_\theta)$ be the configuration reached in ρ^i after reading $u^i(\theta)$. Notice that $K_\theta = \{h_1, \ldots, h_j\}$ where $j = \max\{k \mid \theta_k \leq \theta\}$, and thus does not depend on i. We define similarly $v(\theta)$, q_θ, α_θ, and g_θ for the leader.

We let $n_i(\theta)$ be the sum of the number of capacity reads of h_i occurring in $v(\theta)$, $u^1(\theta), \ldots, u^m(\theta)$.

We claim that $(\sum_{i=1}^m (n_i+1)[p_i], qA, g) \xrightarrow{\overline{w}} (\sum_{i=1}^m [p_i] + n_i[p''_i], qA\alpha', g)$ in the $(\mathcal{C}_{fin}, \mathcal{D})$-system. More precisely, we show by induction on θ that for all $\theta \leq r$, the $(\mathcal{C}_{fin}, \mathcal{D})$-system has a run of the form:

$$(\textstyle\sum_{i=1}^m (n_i+1)[p_i], qA, g) \xrightarrow{\overline{b}_1 \cdots \overline{b}_\theta} (M_\theta, q_\theta\alpha_\theta, g'_\theta)$$

where

$$M_\theta = \textstyle\sum_{\{i \mid \theta < \theta_i\}} (n_i+1)[p^i_\theta] \;+\; \sum_{\{i \mid \theta \geq \theta_i\}} \left([p^i_\theta] + (n_i - n_i(\theta))[p'_i] + n_i(\theta)[p''_i]\right) ,$$

and if b_θ is part of one of the $y^i_{k,j}$, $(y^i_{k,j})^{(n_i+1)}$, or if it is one of the $a_{k,j}$, then $g'_\theta = g_\theta = g^1_\theta = \ldots = g^m_\theta$.

The intuition is that at any time, the main copy of ρ^i is in state p^i_θ. Before the first write of h_i (i.e. $\theta < \theta_i$), the n_i remaining copies progress with the main copy. After the first occurrence of $\overline{w}(h_i)$ (i.e. $\theta \geq \theta_i$), the main copy continues alone. The other copies are either in state p'_i, waiting to write h_i, or stopped in state s''_i after writing h_i. The transition from p'_i to p''_i happens each time the next b_θ corresponds to a capacity read in one of the u^j or v. So after $\overline{b}_1 \cdots \overline{b}_\theta$, there are $n_i(\theta)$ copies in p''_i.

Assume that this holds for some $\theta \geq 0$.

Case 1: $\overline{b}_{\theta+1}$ is the first occurrence of $\overline{w}(h_j)$ for some j, i.e. $\theta + 1 = \theta_j$. Then we have $p^j_\theta = p'_j$, $p^j_{\theta+1} = p''_j$ and $p^i_\theta = p^i_{\theta+1}$ for all $i \neq j$. Moreover, $n_j(\theta) = n_j(\theta+1) = 0$. Thus $M_{\theta+1} = M_\theta - [p'_j] + [p''_j]$, and we can complete

the run of the $(\mathcal{C}_{fin}, \mathcal{D})$-system by letting one of the $n_j + 1$ contributors in state p'_j take the transition $p'_j \xrightarrow{\overline{w}(h_i)} p''_j$.

Case 2: $b_{\theta+1} = \boldsymbol{a_{k,j}}$, i.e $\overline{b}_{\theta+1} = w(h)$, $\overline{b}_{\theta+1} = r(h)$ or $\overline{b}_{\theta+1} = \overline{w}(h)r(h)$. We have $(K_\theta, q_\theta\alpha_\theta, g_\theta) \xrightarrow{a_{k,j}} (K_{\theta+1} = K_\theta, q_{\theta+1}\alpha_{\theta+1}, g_{\theta+1} = h)$ in $\mathcal{D}^\kappa$, and for all i, $(K_\theta, p_\theta, g_\theta) \xrightarrow{a_{k,j}} (K_{\theta+1} = K_\theta, p_\theta = p_{\theta+1}, g_{\theta+1} = h)$ in $\mathcal{C}^\kappa$. So the second property we have to prove, $g_{\theta+1} = g^1_{\theta+1} = \ldots = g^m_{\theta+1}$, is true. We only need to show that $(M_\theta, q_\theta\alpha_\theta, g'_\theta) \xrightarrow{\overline{b}_{\theta+1}} (M_{\theta+1}, q_{\theta+1}\alpha_{\theta+1}, h)$.

- If $\overline{b}_{\theta+1} = w(h)$, this is immediate.
- If $\overline{b}_{\theta+1} = r(h)$, then $r(h)$ is a register read in v (and all of the u^i), that is, $h \notin K_\theta = K_{\theta+1}$, and $g_\theta = g^i_\theta = \ldots = g^m_\theta = h$. For all i, we have $g^i_\theta = last(u^i(\theta)) = last(y^i_{k,1}z^i_{k,1}a_{k,1}\ldots y^i_{k,j}z^i_{k,j})$ (the last equality holds assuming $z^i_{k,j} \neq \varepsilon$). Since $z^i_{k,j}$ only contains writes and reads of values in K_θ, we must have $z^i_{k,j} = \varepsilon$ for all i. Then by induction hypothesis, we have $g'_\theta = g_\theta = g^1_\theta = \ldots = g^m_\theta = h$. Moreover, $M_\theta = M_{\theta+1}$, so we indeed have $(M_\theta, q_\theta\alpha_\theta, h) \xrightarrow{r(h)} (M_{\theta+1}, q_{\theta+1}\alpha_{\theta+1}, h)$.
- If $\overline{b}_{\theta+1} = \overline{w}(h)r(h)$, then $r(h_j)$ is a capacity read in v, and thus must occur after $\nu(h_j)$, i.e. $\theta \geq \theta_j$. We also have $n_j(\theta+1) = n_j(\theta)+1$, thus $M_{\theta+1} = M_\theta - [p'_j] + [p''_j]$. So $(M_\theta, q_\theta\alpha_\theta, g'_\theta) \xrightarrow{\overline{w}(h)} (M_{\theta+1}, q_\theta\alpha_\theta, h) \xrightarrow{r(h)} (M_{\theta+1}, q_{\theta+1}\alpha_{\theta+1}, h)$.

Case 3: $b_{\theta+1}$ is part of one of the $y^i_{k,j}$ or $(y^i_{k,j})^{(n_i+1)}$. Then it corresponds to a register read in one of the u^i, sauy $i = i_0$. Thus, $\overline{b}_{\theta+1} = \overline{r}(h)$ or $\overline{b}_{\theta+1} = \overline{r}(h)^{n_{i_0}+1}$. We first show that $g'_\theta = g_\theta = g^1_\theta = \ldots = g^m_\theta = h$. We have $(K_\theta, p^{i_0}_\theta, g^{i_0}_\theta = h) \xrightarrow{\overline{r}(h)} (K_{\theta+1} = K_\theta, p^{i_0}_{\theta+1}, g^{i_0}_{\theta+1} = h)$, and $h \notin K_\theta$. If $\theta = 0$, then $g^{i_0}_\theta = g'_\theta = g = h$. If $\theta > 0$, by construction of w, b_θ is either also part of some $y^i_{k,j}$ for $i \leq i_0$, or $\boldsymbol{a_{k,j-1}}$, or $\overline{w}(h_{k-1})$ (then $j = 1$). The case $b_\theta = \overline{w}(h_{k-1})$ is in fact impossible: $u^i(\theta)$ would end with $\nu(h_{k-1})$, and we would have $g^{i_0}_\theta = h = h_{k-1} \in K_\theta$, which contradicts $h \notin K_\theta$. So b_θ is either part of some $y^i_{k,j}$, or $\boldsymbol{a_{k,j-1}}$. By induction hypothesis, and since $g^{i_0}_\theta = h$, we obtain $g'_\theta = g_\theta = g^1_\theta = \ldots = g^m_\theta = h$.

Since $g^i_\theta = g^i_{\theta+1}$ for all $i \neq i_0$ and $g_\theta = g_{\theta+1}$, we also have $g_{\theta+1} = g^1_{\theta+1} = \ldots = g^m_{\theta+1} = h$. It remains to show that $(M_\theta, q_\theta\alpha_\theta, h) \xrightarrow{\overline{b}_{\theta+1}} (M_{\theta+1}, q_{\theta+1}\alpha_{\theta+1} = q_\theta\alpha_\theta, h)$, i.e., $M_\theta \xrightarrow{\overline{b}_{\theta+1}} M_{\theta+1}$.

- If $\overline{b}_{\theta+1} = \overline{r}(h)^{(n_{i_0}+1)}$, then $\theta_{i_0} > \theta + 1$, so $M_\theta \geq (n_{i_0}+1)[p^{i_0}_\theta]$ and $M_{\theta+1} = M_\theta - (n_i+1)[p^{i_0}_\theta] + (n_i+1)[p^{i_0}_{\theta+1}]$. Thus $M_\theta \xrightarrow{\overline{r}(h)^{(n_{i_0}+1)}} M_{\theta+1}$.

- If $\overline{b}_{\theta+1} = \overline{r}(h)$, then $\theta_{i_0} < \theta$, and $M_{\theta+1} = M_\theta - [p^{i_0}_\theta] + [p^{i_0}_{\theta+1}]$, so $M_\theta \xrightarrow{\overline{r}(h)} M_{\theta+1}$.

Case 4: $b_{\theta+1}$ is part of one of the $z^i_{k,j}$ or $(z^i_{k,j})^{(n_i+1)}$. This is similar to case 2 (3rd item). □

Proof of Lemma 40. The idea is that contributors starting in p_i behave as in the run of Lemma 39, while the contributors starting in p''_i wait until the main copy of ρ^i reaches p''_i, and then follow it for the part $p''_i \xrightarrow{u^i_{i+1}\cdots u^i_{m+1}} p_i$. We do not give all details of the proof, which are very similar to the proof of Lemma 39. We only explain how to define the run of the $(\mathcal{C},\mathcal{D})$-system, and state the invariant for the induction.

We let

$$\begin{aligned} w'_k = {} & (y^1_{k,1})^{(n_1+1)} \cdots (y^m_{k,1})^{(n_k+1)} (z^1_{k,1})^{(n_k+1)} \cdots (z^m_{k,1})^{(n_m+1)} \boldsymbol{a_{k,1}} \\ & \cdots \\ & (y^1_{k,\ell_k})^{(n_1+1)} \cdots (y^m_{k,\ell_k})^{(n_k+1)} (z^1_{k,\ell_k})^{(n_k+1)} \cdots (z^m_{k,\ell_k})^{(n_m+1)} \boldsymbol{a_{k,\ell_k}} \\ & (y^1_{k,\ell_k+1})^{(n_1+1)} \cdots (y^m_{k,\ell_k+1})^{(n_k+1)} (z^1_{k,\ell_k+1})^{(n_k+1)} \cdots (z^m_{k,\ell_k+1})^{(n_m+1)} \overline{w}(h_k)\,, \end{aligned}$$

and w', b'_θ, $\overline{b}'_\theta$, $\overline{w}$ as before. One can show by induction on θ that the $(\mathcal{C}_{fin},\mathcal{D})$-system has a run of the form:

$$\begin{aligned} & (\textstyle\sum_{i=1}^m (n_i+1)[p_i] + n_i[p''_i], tA, g) \xrightarrow{\overline{w}_1\cdots\overline{w}_\theta} \\ & \left(\begin{array}{l} \sum_{\{i \mid \theta<\theta_i\}} \left((n_i+1)[p^i(\theta)] + n_i[p''_i]\right) + \\ \sum_{\{i \mid \theta\geq\theta_i\}} \left((n_i+1)[p^i(\theta)] + (n_i - n_i(\theta))[p'_i] + n_i(\theta)[p''_i]\right) \end{array} , q_\theta \alpha_\theta, g' \right). \end{aligned}$$

□

C $(\mathcal{C},\mathcal{D})$-systems with multiple registers

The definition of $(\mathcal{C},\mathcal{D})$-systems extends to systems with m registers as expected. Given a set G of register values, we let

$$\begin{aligned} \Sigma^m_C &= \{\overline{r}_i(g), \overline{w}_i(g) : 1 \leq i \leq m, g \in G\}\,, \qquad \text{and} \\ \Sigma^m_D &= \{r_i(g), w_i(g) : 1 \leq i \leq m, g \in G\}\,. \end{aligned}$$

For instance, $\overline{r}_i(g)$ corresponds to a contributor read from the i-th register.

Assume $\mathcal{C}$, $\mathcal{D}$ are transition systems over Σ^m_C and Σ^m_D, resp. Configurations of the $(\mathcal{C},\mathcal{D})$-system are described as tuples

$$(M \in \mathbb{N}^S, t \in T, g_1 \in G, \ldots, g_n \in G)\,.$$

Transitions are defined similarly to the case of $(\mathcal{C},\mathcal{D})$-systems with one register:

$$(M,t,g_1,\dots g_i,\dots,g_n) \xrightarrow{w_i(h)} (M,t',g_1,\dots h,\dots,g_n) \quad \text{if } t \xrightarrow{w_i(h)} t' \text{ in } \Delta\,,$$
$$(M,t,g_1,\dots g_i,\dots,g_n) \xrightarrow{r_i(h)} (M,t',g_1,\dots h,\dots,g_n) \quad \text{if } t \xrightarrow{r_i(h)} t' \text{ in } \Delta \text{ and } h = g_i\,,$$
$$(M,t,g_1,\dots g_i,\dots,g_n) \xrightarrow{\overline{w}_i(h)} (M',t,g_1,\dots h,\dots,g_n) \quad \text{if } M \xrightarrow{\overline{w}_i(h)} M' \text{ in } \delta\,,$$
$$(M,t,g_1,\dots g_i,\dots,g_n) \xrightarrow{\overline{r}_i(h)} (M',t,g_1,\dots h,\dots,g_n) \quad \text{if } M \xrightarrow{\overline{r}_i(h)} M' \text{ in } \delta \text{ and } h = g_i\,.$$

From a $(\mathcal{C},\mathcal{D})$-system with m registers, we can construct an "equivalent" $(\mathcal{C}',\mathcal{D}')$-system with 1 register, using the ideas of Section 7. The contents of the m registers will be stored in the states of $\mathcal{D}'$, and the unique register of the $(\mathcal{C}',\mathcal{D}')$-system will be used to communicate about the actions of the contributors. So the values of the register are:

$$G' = \{\overline{\mathtt{r_i?}}(g), \overline{\mathtt{w_i?}}(g), \overline{\mathtt{r}}_\mathtt{i}(g), \overline{\mathtt{w}}_\mathtt{i}(g), \mathtt{r_i}(g), \mathtt{w_i}(g) \mid 1 \le i \le m \text{ and } g \in G\}$$

and the alphabets of $\mathcal{C}'$ and $\mathcal{D}'$ are defined as usual:

$$\Sigma'_C = \{\overline{r}(a), \overline{w}(a) \mid a \in G'\}, \qquad \Sigma'_D = \{r(a), w(a) \mid a \in G'\}\,.$$

For a word $u \in (\Sigma'_C \cup \Sigma'_D)^\infty$, we define as before $\mathit{trans}(u)$ by removing all actions of the contributors and all reads of the leader, and replacing leader writes $w(a)$ by a.

The states of $\mathcal{D}'$ are:

$$T' = T \times G^m \cup \{[t, x_1 \dots, x_n] \mid x_i = \overline{\mathtt{r}}_\mathtt{i}(g) \text{ or } x_i = \overline{\mathtt{w}}_\mathtt{i}(g) \text{ for some } i, \text{ and } x_i \in G \text{ for all } j \ne i\}$$

and the states of $\mathcal{C}'$:

$$S' = S \cup \{[s,a,s'] \mid s,s' \in S \text{ and } a = \overline{\mathtt{r_i?}}(g) \text{ or } a = \overline{\mathtt{w_i?}}(g) \text{ for some } i,g\}\,.$$

The transitions of $\mathcal{C}'$ are defined as in Section 7:

$$s \xrightarrow{\overline{w}(\overline{\mathtt{w_i?}}(g))} [s, \overline{\mathtt{w}}_\mathtt{i}(g), s'] \xrightarrow{\overline{r}(\overline{\mathtt{w}}_\mathtt{i}(g))} s' \quad \text{if} \quad s \xrightarrow{\overline{w}_i(g)} s' \text{ in } \mathcal{C}$$
$$s \xrightarrow{\overline{w}(\overline{\mathtt{r?}}(g))} [s, \overline{\mathtt{r}}_\mathtt{i}(g), s'] \xrightarrow{\overline{r}(\overline{\mathtt{r}}_\mathtt{i}(g))} s' \quad \text{if} \quad s \xrightarrow{\overline{r}_i(g)} s' \text{ in } \mathcal{C}$$

and similarly for transitions of the leader, except $\mathcal{D}'$ now keeps the values of

m registers:

$$[t, g_1, \ldots, g_i, \ldots, g_n] \xrightarrow{w(\mathtt{w_i}(h))} [t', g_1, \ldots, h, \ldots, g_n] \quad \text{if} \quad t \xrightarrow{w_i(h)} t' \text{ in } \mathcal{D}$$

$$[t, g_1, \ldots, g_i, \ldots, g_n] \xrightarrow{w(\mathtt{r}(g_i))} [t', g_1, \ldots, g_i, \ldots, g_n] \quad \text{if} \quad t \xrightarrow{r_i(g_i)} t' \text{ in } \mathcal{D}$$

$$[t, g_1, \ldots, g_i, \ldots, g_n] \xrightarrow{r(\overline{\mathtt{w_i?}}(h))} [t, g_1, \ldots, \overline{\mathtt{w}}_\mathtt{i}(h), \ldots, g_n] \xrightarrow{w(\overline{\mathtt{w}}_\mathtt{i}(h))} [t, g_1, \ldots, h, \ldots, g_n]$$

$$[t, g_1, \ldots, g_i, \ldots, g_n] \xrightarrow{r(\overline{\mathtt{r_i?}}(g_i))} [t, g_1, \ldots, \overline{\mathtt{r}}_\mathtt{i}(g_i), \ldots, g_n] \xrightarrow{w(\overline{\mathtt{r}}_\mathtt{i}(g_i))} [t, g_1, \ldots, g_i, \ldots, g_n]\,.$$

Theorem 41 *1. For every trace u of the $(\mathcal{C}, \mathcal{D})$-system, there exists a trace u' of the $(\mathcal{C}', \mathcal{D}')$-system such that $trans(u') = u$.*

If u is finite and and the $(\mathcal{C}, \mathcal{D})$-system has a run over u ending in $(M, t, g_1, \ldots, g_n)$, then the $(\mathcal{C}', \mathcal{D}')$-system has a run over u' ending in $(M, [t, g_1, \ldots, g_n], a)$ where a is the last action of u.

2. For every trace u' of the $(\mathcal{C}', \mathcal{D}')$-system, there exists a trace $u \in stutt(trans(u'))$ in the $(\mathcal{C}, \mathcal{D})$-system.

If u' is finite and the $(\mathcal{C}', \mathcal{D}')$-system has a run over u' ending in $(M, [t, g_1, \ldots, g_n], a)$ with $M \in \mathbb{N}^S$ and $g_1, \ldots, g_n \in G$, then the $(\mathcal{C}, \mathcal{D})$-system has a run over u ending in $(M, t, g_1, \ldots, g_n)$.

As in Section 7, we can also modify the $(\mathcal{C}', \mathcal{D}')$-system to preserve maximality of runs, and prove Theorem 38.